# WATERING THE NEW ECONOMY

## Managing the impacts of the AI revolution

January 2026

GLOBAL WATER INTELLIGENCE

# FOREWORD

**Matthew Pine**
President and Chief Executive Officer, Xylem

*The Artificial Intelligence supercycle is accelerating faster than any industrial shift in history. It is reshaping economies and redefining competitiveness. But this AI revolution has a hidden vulnerability: every chip, every data center, every megawatt of power depends on reliable supplies of clean water.*

Until recently, debates about AI focused on familiar constraints like energy and raw materials. Now industry, communities, and policymakers are recognizing AI's impact on a world that's already water stressed. The search for solutions, however, has been hampered by a lack of data and understanding of AI. Watering the New Economy offers a strategic, fact-based framework for understanding both the challenges and opportunities.

This report maps AI's pressure points on our world's water supplies, which could pit industries against communities as they compete for a shrinking resource. Most importantly, it highlights pathways that can enable us to achieve both goals: to protect affordable, healthy water for communities and safeguard the innovation benefits of the AI economy.

The choice is clear. Managed poorly, demand for water could turn into a zero-sum contest between people and progress. Managed well, it can become a catalyst for something larger: a global transition to greater water security.

Our goal was to give all decision makers – from policymakers to utilities, tech companies, and investors – a shared framework for a broad dialogue about how we can solve water for the AI economy. It's time for a broad-based water transition. The moment to mobilize technology, capital, and collaboration is now.

# Contents

# Executive summary

## Turning a threat into an opportunity

The AI revolution is reshaping the global economy, and with it, the way we manage water. Embracing the benefits of AI requires accommodating a material increase in water demand across the value chain. However, the new economy, water utilities, and communities can succeed through partnerships to expand water infrastructure and supply at manageable costs. These partnerships, centering on efficiency, circularity, and infrastructure renewal, form the crux of a Water Transition for sustainable growth.

### The three industries at the heart of the new economy for AI are data centers, chip fabs and power generation.

This new economy is a major and growing user of water. In 2025, it withdrew 23.7km³, a 38% increase over 2020. Water use will grow another 129% by 2050. Even so, it is a far less water-intensive sector than other major drivers of the economy – in 2025 it represented just 3.7% of total industrial water withdrawals. While AI's water footprint is not as intensive as other sectors, it is attracting attention for four reasons:

**1 Growth**
The rapid pace of AI adoption is loading additional demand onto existing infrastructure at speed and scale.

**2 Location**
Water rarely drives siting decisions, so semiconductor fabs and data centers continue to cluster in water-stressed regions.

**3 Timing**
Rising demand comes as climate extremes make the water cycle more volatile and less reliable.

**4 System weakness**
Aging public water infrastructure is unprepared after decades of underinvestment.

AI's water footprint is not where most people think it is. While public attention often focuses on data center cooling, this is actually the least water-intensive piece of the AI value chain, which also includes the water used to make semiconductor chips and produce the power for data centers and fabs. Three key levers will help society accommodate the impact of AI-driven growth:

**Network efficiency and reuse:** Leakage reduction and integrated reuse can neutralize withdrawal growth, keeping future demand at today's levels.

**Boosting onsite efficiency:** Applying best practices from leading chip fabs and data centers to the rest of the sector.

**Optimising the energy mix:** Supporting renewable energy integration and deployment, and shifts from coal to gas.

### How can we limit future withdrawals to today's levels?

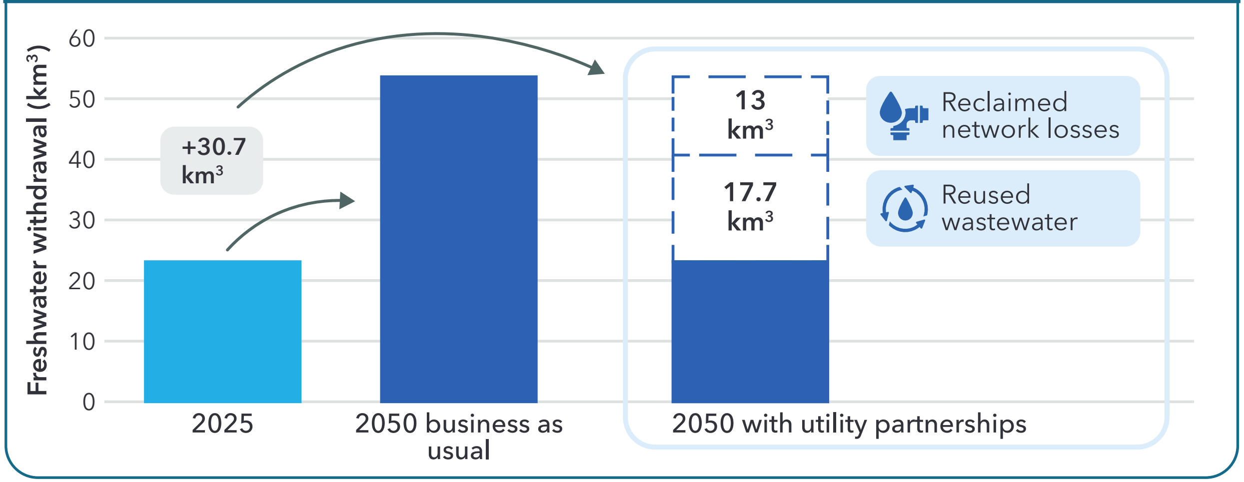



Cities need reliable, resilient water systems, while new economy firms need supply security and are prepared to invest. Cross-sector partnerships can:

- Deliver integrated reuse and leakage reduction projects that serve both municipal and industrial needs.
- Launch a 'Water Transition' that builds climate resilience and unlocks AI-driven economic growth.

The challenge is clear: we need to act quickly to allow water to be a catalyst of the new economy.

# Water in economic transitions

All economic transitions demand changes in water management. This one is no different

While the water risks of the AI revolution are not as externalized as past pollution, they still present a threat as an additional draw on shared – and sometimes already stretched – water resources. As the economy evolves, so must the ways we manage and value water.

**Water drives industry**

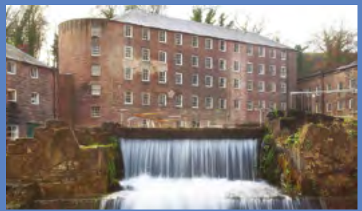

**Industry transforms rivers**

**Overuse ignites reform**

**Innovation tests new limits**

From the textile mills of the 1800s to today's AI data centers, every major economic shift has redrawn the map of water use. Heavy industry originally concentrated on major rivers and lakes until pollution spurred regulation. The new economy's industries use less water than steel or chemicals ever did, but they are emerging in a world where water systems are already stressed. Rising global temperatures are intensifying droughts and floods. At the same time, years of underinvestment have left public water infrastructure extremely vulnerable.

Adjusting to new patterns of water use in industry has always been painful. It took the old fossil-fuel-powered manufacturing economy more than 100 years to reckon with its water impact, with the passage of the Clean Water Act in the early 1970s. The good news is that the data center and chip industries are willing to engage right now. The stability of their operations are at stake. This whitepaper outlines a strategy for turning that economic risk into a resilient future for us all.

## Water management and industrial geography

### Early industry followed water

Textile mills of the 1800s clustered in valleys with a strong head of water, using rivers for both power and process needs. Conflicts between users drove the creation of some of the first shared water management systems. By 1900, industry accounted for 10–15% of global freshwater use.

### Heavy industry shaped modern cities

Steel, chemical, and paper manufacturing – major components of the 'old economy' – concentrated along major rivers and lakes to secure vast supplies for cooling and discharge. Over time, these waters became polluted, prompting new governance models and regulation.

### AI is breaking the old pattern

Traditional cloud data centers located near large populations to minimize latency. Data centers used for training and running AI models have different priorities: low cost power and real estate are more important. Water availability is not a significant factor in determining location, creating operational and strategic risks.

### A new challenge for water managers

Although AI's water use is modest compared to legacy industries, its decoupling from water-rich regions demands proactive strategies for sustainable supply, recycling, and local resilience.

# AI's water footprint

## AI uses water in three ways: onsite data center cooling, offsite power generation, and chip fabrication

This new analysis shows that 30 minutes of AI use requires just over 600ml of water, much less than previous estimates of the per-query water use of AI. The biggest impact is through the power supply to data centers, followed by the water embedded in chip manufacturing and onsite data center cooling requirements.

Within the new economy, AI is the most significant engine of growth. Its deployment relies on energy-intensive data centers and semiconductor fabs, with the latter supplying the chips and servers essential for AI and cloud infrastructure.

### The water impact of 30 minutes of AI use (ml)

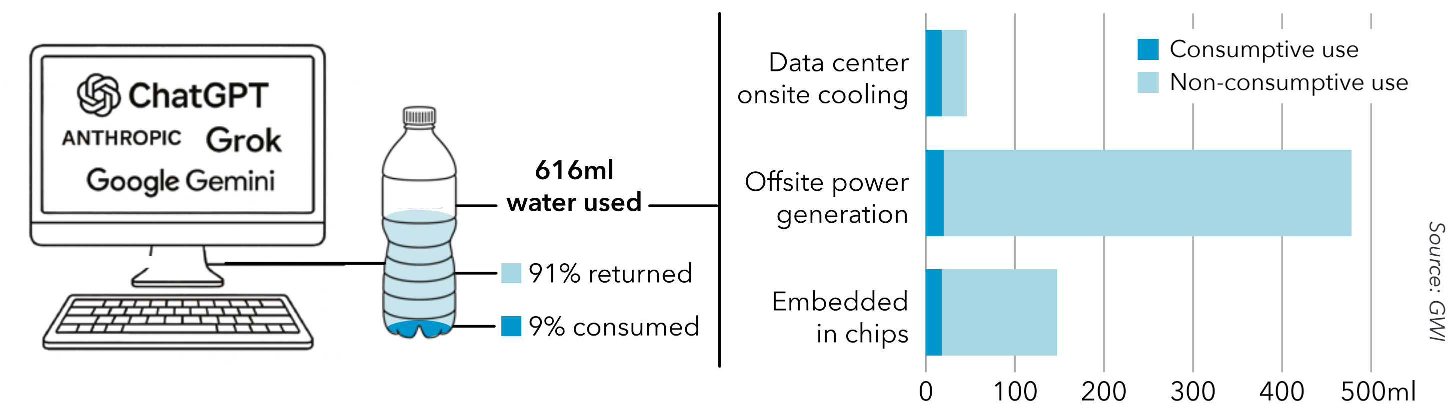


### Usage versus consumption: both matter

Understanding the full water impact of industry requires distinguishing between consumptive and non-consumptive use. Water withdrawn but returned – often warmer or altered – can still degrade ecosystems and constrain other users. Consumptive losses, meanwhile, permanently remove water from the basin. In the new economy, this is primarily through evaporation of cooling water. Both affect availability, quality, and resilience, demanding integrated management and transparency.

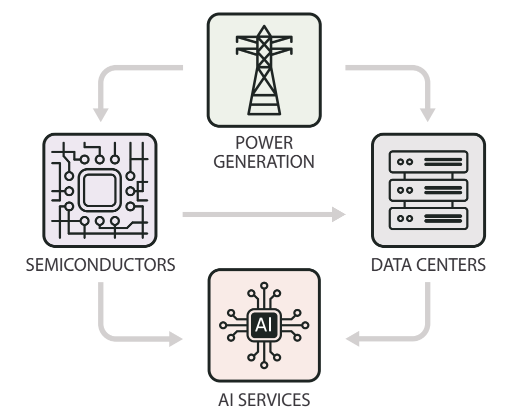


## Methodology

### Data center cooling

Based on analysis of the water use effectiveness of different cooling technologies and their current application in the industry.

### Power generation

Uses the reported power supply mix of major data center operators and applies standard water use intensities for coal, gas and nuclear to the thermoelectric portion. Power use per 30 minutes is 32.5Wh, based on estimates of intensive AI usage.

### Chips

Uses reported water consumption in manufacturing to calculate water embedded in GPUs divides this by 2.5 year life cycle and processing rates.

### Operational vs manufacturing footprints

The chip-embedded footprint is relatively small compared to the semiconductor sector's role in total new economy water use. This is because the per-30 minute figure reflects a single chip's lifecycle, whereas the industry as a whole produces chips both for new data centers and for retrofit upgrades. This embedded water accounts only for the processing chip in a data center, not the user's own device.

# Water and the new economy

## It is not only data centers that need water: chip manufacturing and renewable energy will also shape future demand

The transition to a digital and decarbonized economy is reshaping global water use and driving significant demand in the new economy industries. Water use by data centers, semiconductor fabrication plants, and associated power generation is expected to rise 129% by 2050.

These three sectors form the backbone of the new economy, yet all depend heavily on reliable water supplies for cooling and processing. Water use dynamics are shifting across the new economy as a result of environmental pressures and process changes. Water intensity, the amount of water used per unit of output, is rising in some areas and falling in others.

### New water impacts

| | |
|---|---|
| **Data centers:** although not all data centers use water for cooling, and water use efficiency is improving in those that do, the surge in demand for AI and cloud services means total water use will still rise. | **Assumptions:** data centers will become 46% more efficient in their water use by 2050 as a result of cooling system changes. |
| **Semiconductor fabs:** these facilities rely heavily on ultrapure water to produce chips. Recycling is a priority, but recirculation can concentrate difficult-to-detect low-molecular-weight organics, raising treatment concerns. The sector will see the fastest water demand growth without new investment to improve reuse. | **Assumptions:** semiconductor manufacturing will become 113% more water intensive by 2050, but growing reuse will offset this by 27%. |
| **Power generation:** water use here is mostly a function of the energy mix. The sector is transitioning from once-through cooling systems and coal-fired plants (which use more water per kWh) to renewables and combined-cycle gas. | **Assumptions:** overall water intensity in power generation is projected to fall 82% by 2050 as a result of these power mix changes. |

### New water risks

The risk lies not in water volume but in when and where it's needed. Water is fast becoming a material risk for the new economy, shaping both resource security and social license to operate.

Multiple factors will shape the new economy's growth and water intensity, either amplifying or reducing its impact. Our scenario projects water demand to more than double by 2050, exceeding 54 km³ annually, equal to California's surface water storage capacity or nearly 22 million Olympic swimming pools.

**Forecast water demand in new economy industries**

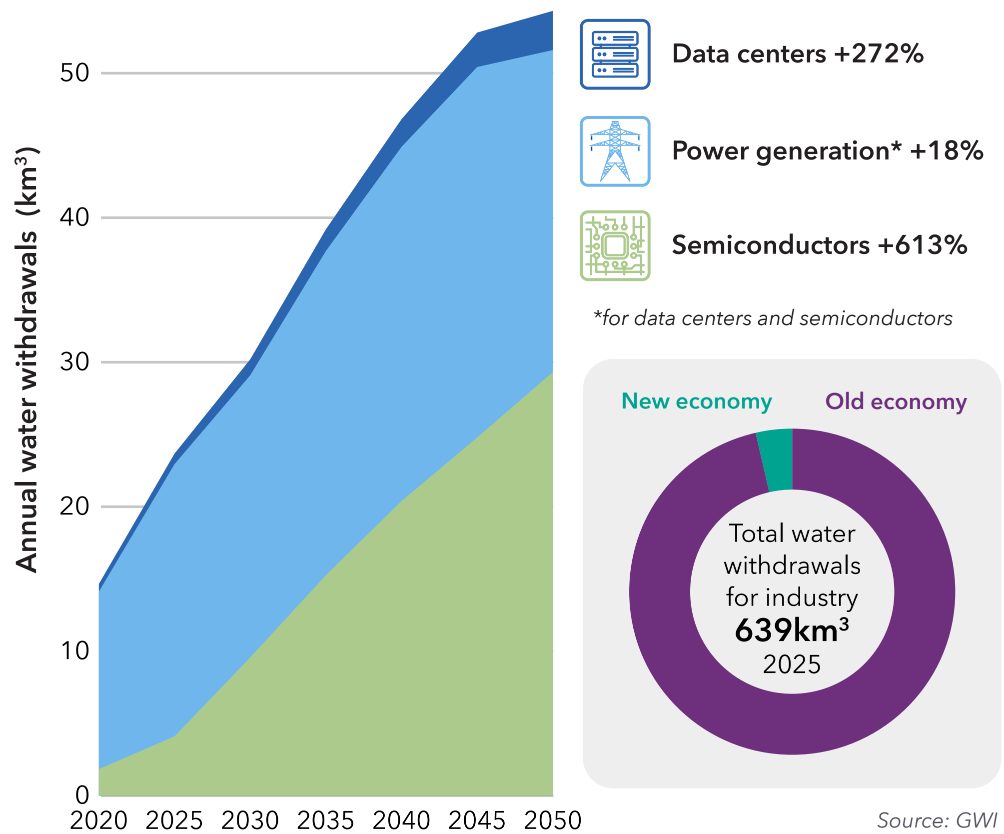


Source: GWI

# Semiconductor water risks

## Chip fabs face rising water intensity and continued exposure to water scarcity

Semiconductor water demand could grow over 600% by 2050, fueled by rapid data center demand and rising per-chip water intensity. Much of this growth will be in water-scarce regions.

Due to the microscopic scale at which chips are produced, water needed during the manufacturing process must be of ultrapure quality. Demand for this process water is expected to rise considerably with the manufacturing demands of AI-specialized chips. Additionally, new fabs continue to cluster in a few key regions, primarily Taiwan and the southwest of the United States, areas with limited water resources.

### Water use in chip manufacturing

Making ultrapure water is highly water intensive. Up to 4m³ of feedwater is required to produce 1m³ of ultrapure water. As chips become more advanced, rising purity standards are further increasing the water intensity of its production.

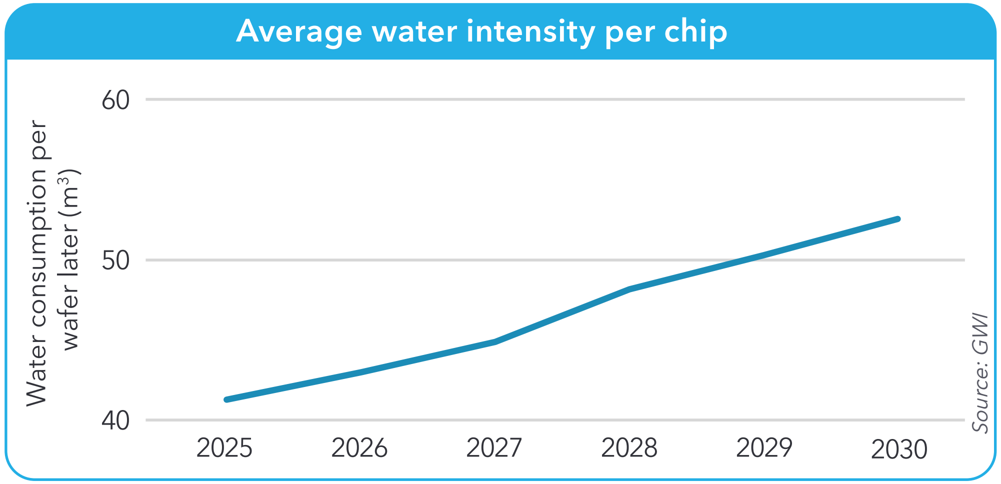


Explosive growth in global demand for semiconductors is being matched by rising water intensity in chip fabrication. Each new generation of manufacturing adds more etching and cleaning steps per wafer, sharply increasing the need for ultrapure water. As a result, semiconductor manufacturing water use intensity will rise 27% by 2030. Yet fabs remain concentrated in water-stressed regions where water availability is already limited. Because relocation is unlikely, large-scale water reuse and recycling will be critical to sustaining both chip production and local water security.

### Advanced fab locations

Manufacturers prioritize infrastructure access, regulatory incentives and low risk of natural disasters for fab siting. Water availability is not a deciding factor, meaning that they are often located in areas of high water stress.

Fears over supply chain vulnerability and the sector's concentration in east Asia has triggered reshoring efforts in the US, with significant funding for fab developers to expand operations in the US. New developments are under construction in Arizona, Texas, and New York.

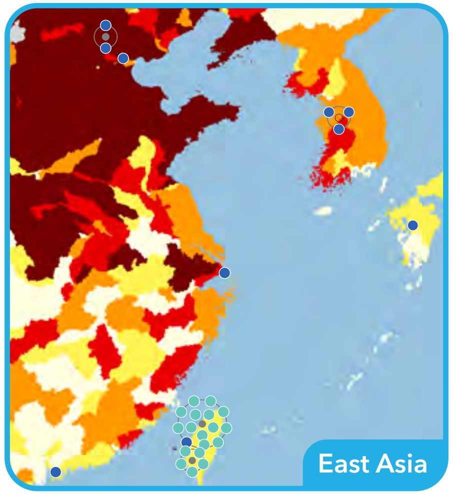


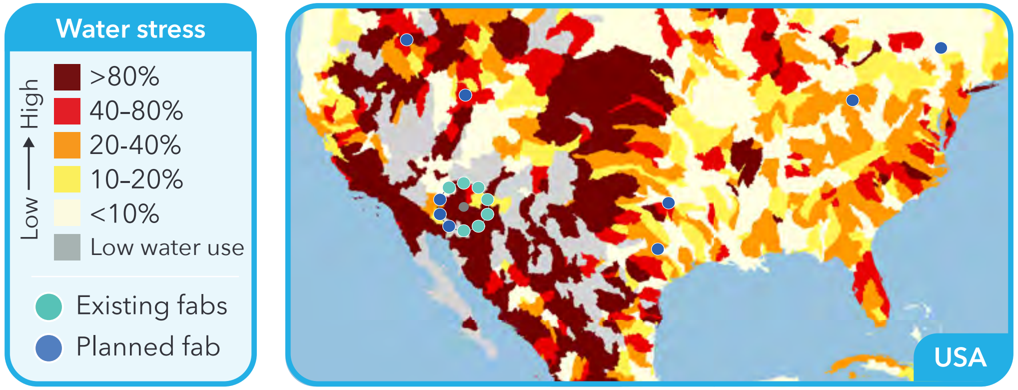


Source: WRI; GWI

# Semiconductor water strategies

## Improved reuse solutions will be a critical enabler of greater supply chain resilience

Rapid industry growth and continued clustering in high-risk regions are intensifying the sector's exposure to scarcity-related water risks. This is necessitating new approaches to water management to mitigate the risks of rising withdrawals.

Semiconductor fabs have long been sited in areas of water stress. But resource pressures are becoming increasingly severe. In Taiwan, drought frequency is increasing, and in Arizona, rising municipal and industrial demand is heightening competition for the sources on which fabs rely. While onsite reuse and recycling are widely deployed, withdrawals in the sector have jumped in recent years. These risks demand renewed focus on water resilience.

### Scarcity and the water challenge

While semiconductor companies are aware of the water risks they currently face, the increasing frequency of extreme weather events is shifting the economics of resilience in the sector.

- During a 2021 drought, the Taiwan Semiconductor Manufacturing Company (the world's largest chipmaker) had to truck water to its Taiwan fabs, costing $28.6 million and driving water expenses up to 2% of quarterly revenue, indicating the financial and operational risks of poor water resilience.
- Such cases reflect the limits of existing strategies and the need for enhanced supply dependability.

### The water strategies available

The sector has deployed advanced wastewater treatment technologies and achieved high reuse rates, yet it rarely recycles ultrapure water back into manufacturing processes. This leaves substantial room for improving circularity and reducing freshwater withdrawal. Cost and a lack of precedent are currently hindering such reuse. But as industrial water tariffs continue to rise in many fab hotspots and water scarcity increasingly shapes business decisions in the sector, manufacturers can adopt full ultrapure water recycling to safeguard water security.

Semiconductor fabs typically repurpose spent ultrapure water for cooling tower make-up. However, onsite reuse in semiconductor fabs is complicated and made more costly by the buildup of low-molecular-weight organics in process wastewater. This limits the extent of water recirculation, and makes some form of discharge or blowdown more operationally attractive. As withdrawals in the semiconductor industry continue to grow from an already high base, fab wastewater can be scaled into a valuable resource for municipal and agricultural reuse.

**Facility water stress of major semiconductor manufacturers**

29% of global semiconductor fabs are in extremely water- stressed areas.

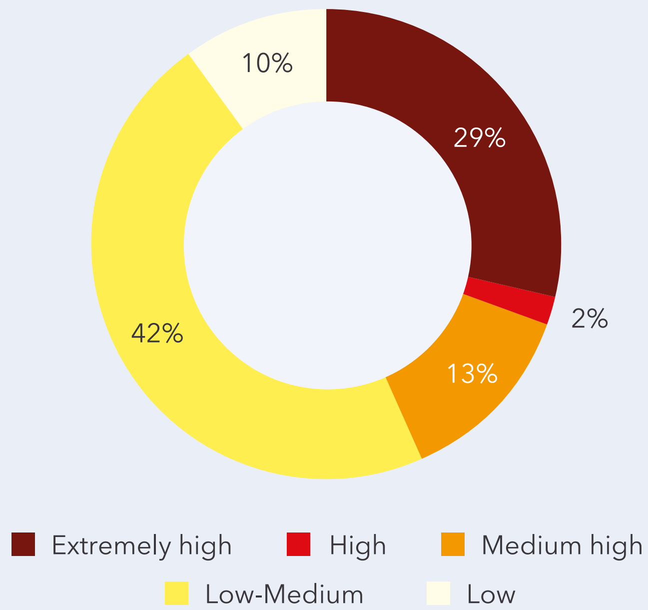


*Source: GWI*

# Water use in data centers

## Data center water efficiency varies considerably based on cooling technology and location

Climate and cooling system choices, which are closely interlinked, determine water efficiency in data centers. Operators deploy diverse cooling strategies, each carrying unique trade-offs for energy use, water consumption, and overall sustainability.

Data centers rely on processing units that consume vast and growing amounts of energy, especially in AI facilities. As computing intensity rises, overheating risks increase, making effective cooling vital. Water's high heat capacity and evaporation properties make it an ideal cooling medium. The easiest way to cool a data center is to open it up and allow air to circulate around it. However this is only possible in certain locations and in certain seasons. Globally, 56% of data center capacity uses some form of evaporative cooling. Within this, cooling choices are determined in large part by environmental conditions. There are three main cooling strategies:

**Cooling towers**, which are more effective in humid climates than other evaporative systems, but are typically less water efficient. In continuous use they typically require around 2.75L/kWh of water to operate, of which **1.85L/kWh** is consumptive use. Where they are used, they are typically used all year round.

**Adiabatic cooling systems**, where air passes through a wetted media before passing into the data center, are increasingly dominant in the sector. They are highly water efficient but rely on more intensive water treatment and cannot be used in humid environments. In continuous operations they typically require around 0.75L/kWh of water, of which **0.63L/kWh** is consumptive use. Usually however adiabatic systems are used between 5-20% of the year, so their average annual water use effectiveness (WUE) is often **below 0.1L/kWh.**

**Direct liquid cooling**, which brings fluid directly to the chip via a closed-loop system, with heat subsequently rejected by a secondary cooling system, such as an air chiller or cooling tower. In continuous operation a direct liquid cooling system with a cooling tower might need around 1.73L/kWh, of which **1.17L/kWh** would be consumptive use. Taking seasonality into account, the average WUE is often **below 0.3L/kWh.**

Direct liquid cooling is an emerging technology necessitated by the high compute density of AI workloads. Higher chip temperature tolerances mean facility-level cooling loops can often rely on air, despite greater heat densities.

### Why don't data center operators use dry cooling methods?

In cool locations, outside air can provide "free cooling," which is useful but geographically limited. Elsewhere, air-cooled chillers are an option, but they shift the burden from water to power. Critically, these systems are needed most on the hottest days, when power grids are already under peak stress. Water-based cooling can be better aligned with system constraints, as water can be stored and managed cost-effectively to meet short periods of peak demand, whereas electricity must be generated in real time and capacity sized for infrequent extremes.

### The seasonality of water use in North American data centers

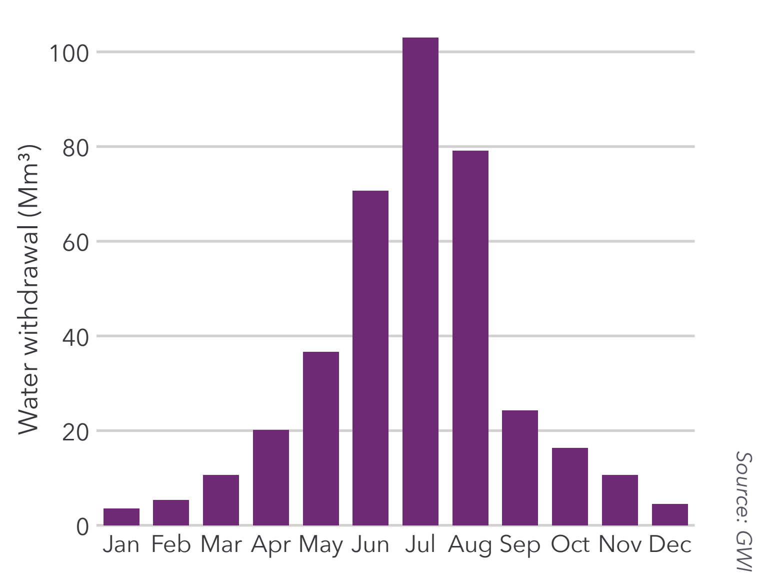


### Seasonality: Less water usage but more water management?

Many data centers only require cooling seasonally, so their total annual water use is far lower than their cooling systems' water use effectiveness (WUE) might suggest. However, this offers little relief from a water management perspective. Water infrastructure must be built to meet peak demand – usually during hot, dry months – so the seasonality of data center cooling actually amplifies pressure on water systems, rather than reducing it.

# Data center water use trends

## The evolution of cooling technology is moving the industry towards greater efficiency, but total water use will still rise

The data center sector is shifting its cooling strategies to manage rising demand and compute intensity sustainably. As a result, data centers are expected to become much more efficient in their water use, limiting the extent of withdrawal growth even as power demand booms.

AI-specialized facilities are scaling direct liquid cooling at the chip level due to the high heat density of advanced chips. Cloud facilities, meanwhile, are shifting to seasonally evaporative adiabatic systems for hybrid facility cooling to curb mounting pressures around water use.

### Trends in cooling

The AI revolution is driving rapid growth of the data center sector, and a diversification of the technologies used for cooling. Conventional air and systems will continue to serve a large installed base, but most new capacity will adopt direct liquid cooling at the chip level to handle higher rack densities. Within these new deployments, approaches vary: some pair direct liquid cooling with secondary water-based systems such as adiabatic or tower/chiller loops, while others rely solely on air for heat rejection. Tower/chiller systems are losing market share as a result of their larger footprint and lower efficiency.

### Water use impacts

Water use in data centers will rise much more slowly than power use – less than half the rate of the nearly seven-fold increase in energy demand expected by 2050. Still, total water withdrawal for cooling will more than triple over the next 25 years, putting pressure on water systems. Cooling innovations mean that sector-wide WUE is improving and expected to drop sharply by 2050. But global averages mask wide variations and the zero-water use of fully air-cooled facilities.

**Cooling system trends 2025-2030**

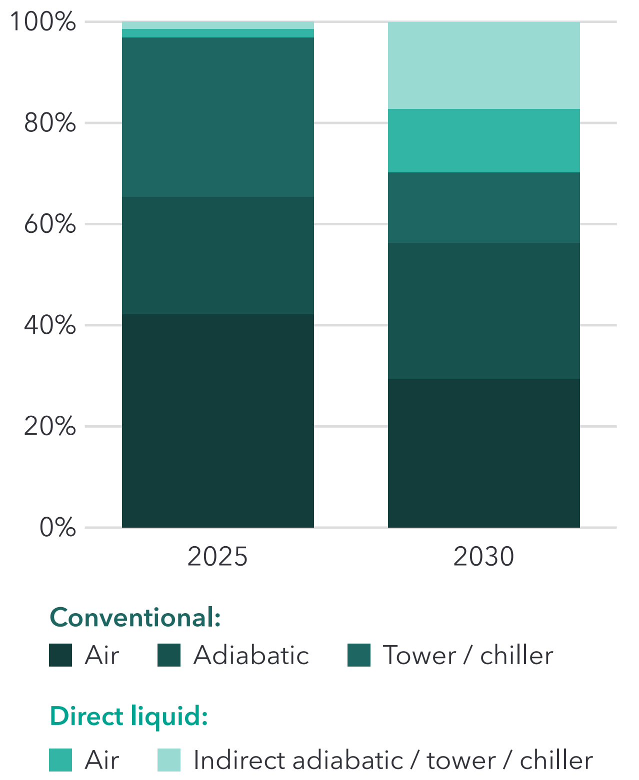


*Source: GWI*

**Forecast data center power and water use 2025-2050 (km³/yr)**

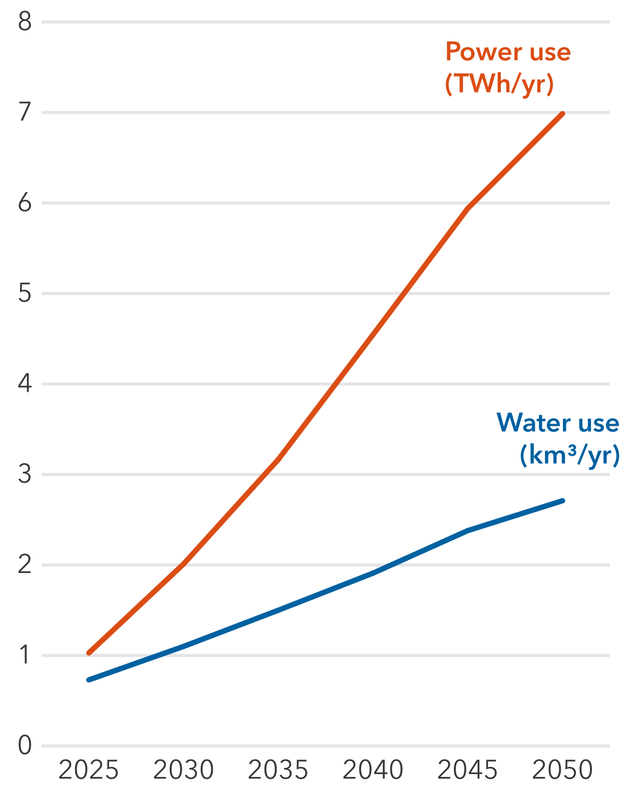


*Source: GWI, S&P Global*

# AI and data center water use in context

Data center water use is significant but it needs to be seen in the context of broader industrial water use

AI uses less water than other aspects of everyday life, and the new economy accounts for only 3.7% of total industrial withdrawals. Nevertheless, it has significant local impacts.

All human activity directly or indirectly uses water. Agricultural products use the most: 190 liters of water are needed to make a cup of coffee with milk. The water embedded in a smart phone is 12,000 liters which spread over their average usage and lifetimes means that they "drink" around 73ml per minute. Desktop computers have a greater mass of manufactured parts and, consequently, a larger water footprint.

## Data center water use in context

A large hyperscale (130MW) data center with a seasonal adiabatic cooling system would use around 171,000 m³ of water in a year. This is enough to supply around 900 households in the US, but it is not much compared to a power plant or a refinery. Some data centers may however use much more water based on their choice of cooling technologies. If all the power supply to the data center came from a coal-fired plant, it would need to intake 79 million m³ in a year to meet the demand. That is enough water to serve a city of 600,000 people.

**Water use in data centers versus other industries ($m^3$/$)**

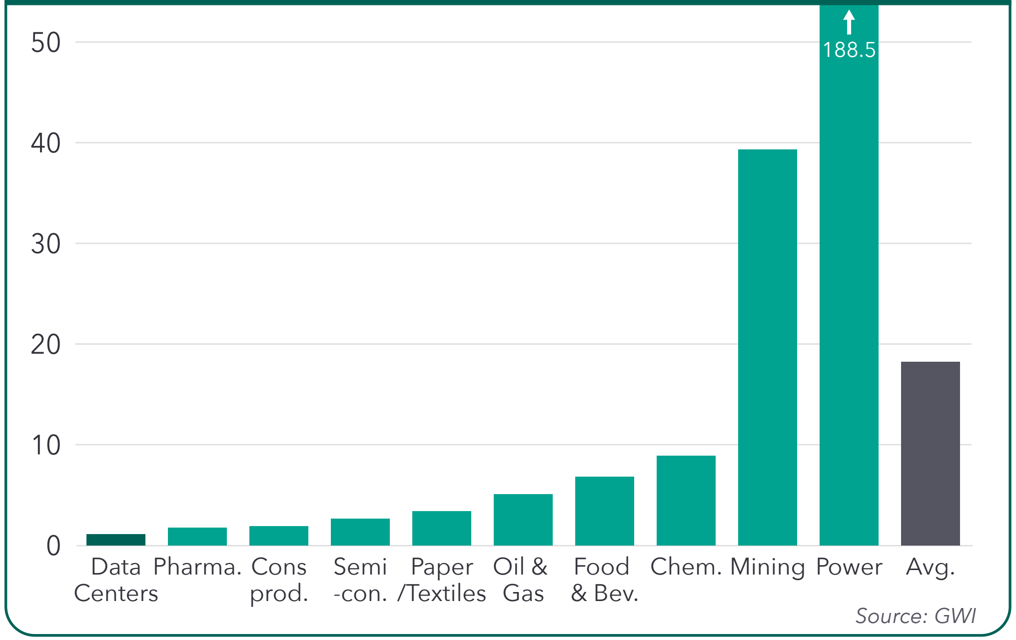


*Source: GWI*

### Water use in daily life


**1 Cup of coffee: 190L**
Direct usage: 0.25L
Power: 0.065L
Cooling: 0L
Embedded: 190L
85% out of basin

**2 30 mins phone browsing: 2.28l**
Direct usage: 0L
Power: 0.062L
Cooling (data centers): 0.02L
Embedded (in device): ~2.2L
9% out of basin

**3 30 mins general activity: 63L**
Direct usage: 0.052L
Power: 3.13L
Cooling: N/A
Embedded: ~59L
63% out of basin

**4 30 mins of AI use: 7.30L**
Direct usage: 0L
Power: 0.458L
Cooling (data centers): 0.028L
Embedded (in personal device): ~7L
10% out of basin

*Source: GWI*

## So why should we be concerned about their water use?

- Data centers are increasingly located in dry or rural areas which have relatively limited water resources.
- Their needs peak in summer when water may be most scarce.
- Much data center use is consumptive: the most efficient cooling systems discharge just 15% of what they withdraw.

These concerns explain why the major players (Amazon, Alphabet, Meta, Microsoft) are committed to being "water positive" by 2030.

# Hyperscale impacts

## Data center water use efficiency also varies by operator based on strategy

Water use effectiveness varies widely among hyperscale data center operators, driven by contrasting strategies for balancing water and energy. Similarly, reporting of water use differs considerably between operators.

All of the major hyperscale operators give figures for their water use effectiveness, although the basis on which this is calculated in not consistent: some report it on the basis of their consumptive use, and some report it on the basis of their total freshwater withdrawals. Not all of them give volumetric data.

### Water strategies

Among the major cloud operators, Amazon, Microsoft, and Meta stand out for ambitious water goals, and all of the players listed here except Alibaba Cloud have committed to replenishing more water than they withdraw by 2030. Operators are also pivoting to alternative and non-potable water sources as a primary strategy in resilience and sustainability. Amazon's data centers are air-cooled for all but the hottest months, and it has pioneered the use of recycled water in direct evaporative cooling. Microsoft has launched a next-generation data center design that consumes no water for cooling in some installations, and in Washington, it partnered with the city of Quincy to build out a dedicated "Water Reuse Utility" that treats cooling blowdown for reuse. Alongside pivoting to recycled sources, operators are also installing onsite treatment and storage capacity for rainwater harvesting.

### Water vs energy

There is a direct trade-off in some circumstances between energy use and water use: using more energy can reduce water use, and vice versa. For instance, mechanical air chillers use a lot of power but no water. Because water is generally underpriced by public-sector providers while electricity is fully priced by private-sector providers, the most cost-effective strategy for some operators is to prioritize conserving energy rather than water. This may explain the wide variation in WUE performance between operators.

**The water use effectiveness of the big five**

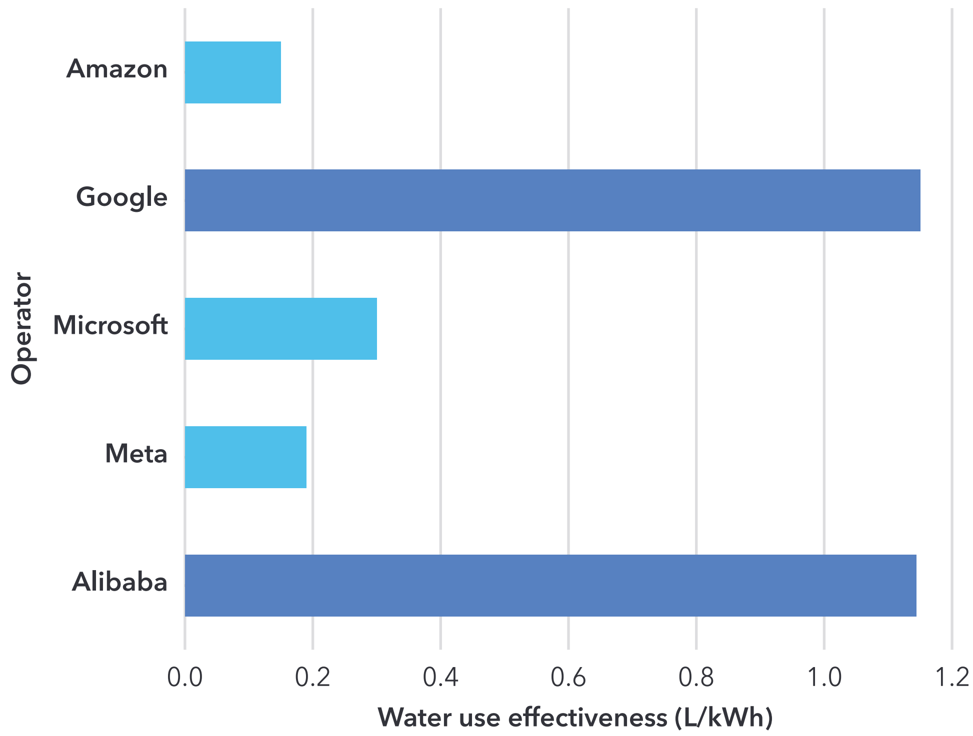


| Operator | Total consumption | Total withdrawal |
|---|---|---|
| Amazon | Not disclosed | Not disclosed |
| Google | 29.48Mm³ | 37.35Mm³ |
| Microsoft | 5.81Mm³ | 10.38Mm³ |
| Meta | 2.97Mm³ | 4.15Mm³ |
| Alibaba | Not disclosed | Not disclosed |

Reported 2024 WUE based on: Total consumptive use | Total withdrawals

*Source: Company sustainability reports and other disclosures; GWI*

# Data center geography

## Data center growth is expected to take place in regions already facing water stress and future water risk

40% of the world's data centers are currently in areas of high or extremely high water stress. The energy demands of the next generation of AI data centers will change the industrial geography of the sector.

Data centers are currently largely located to give good connectivity with population centers in high- and upper-middle-income countries. Growing energy demands are forcing new data center developments to prioritize energy access in site selection.

### Where are the biggest data center clusters in the world?

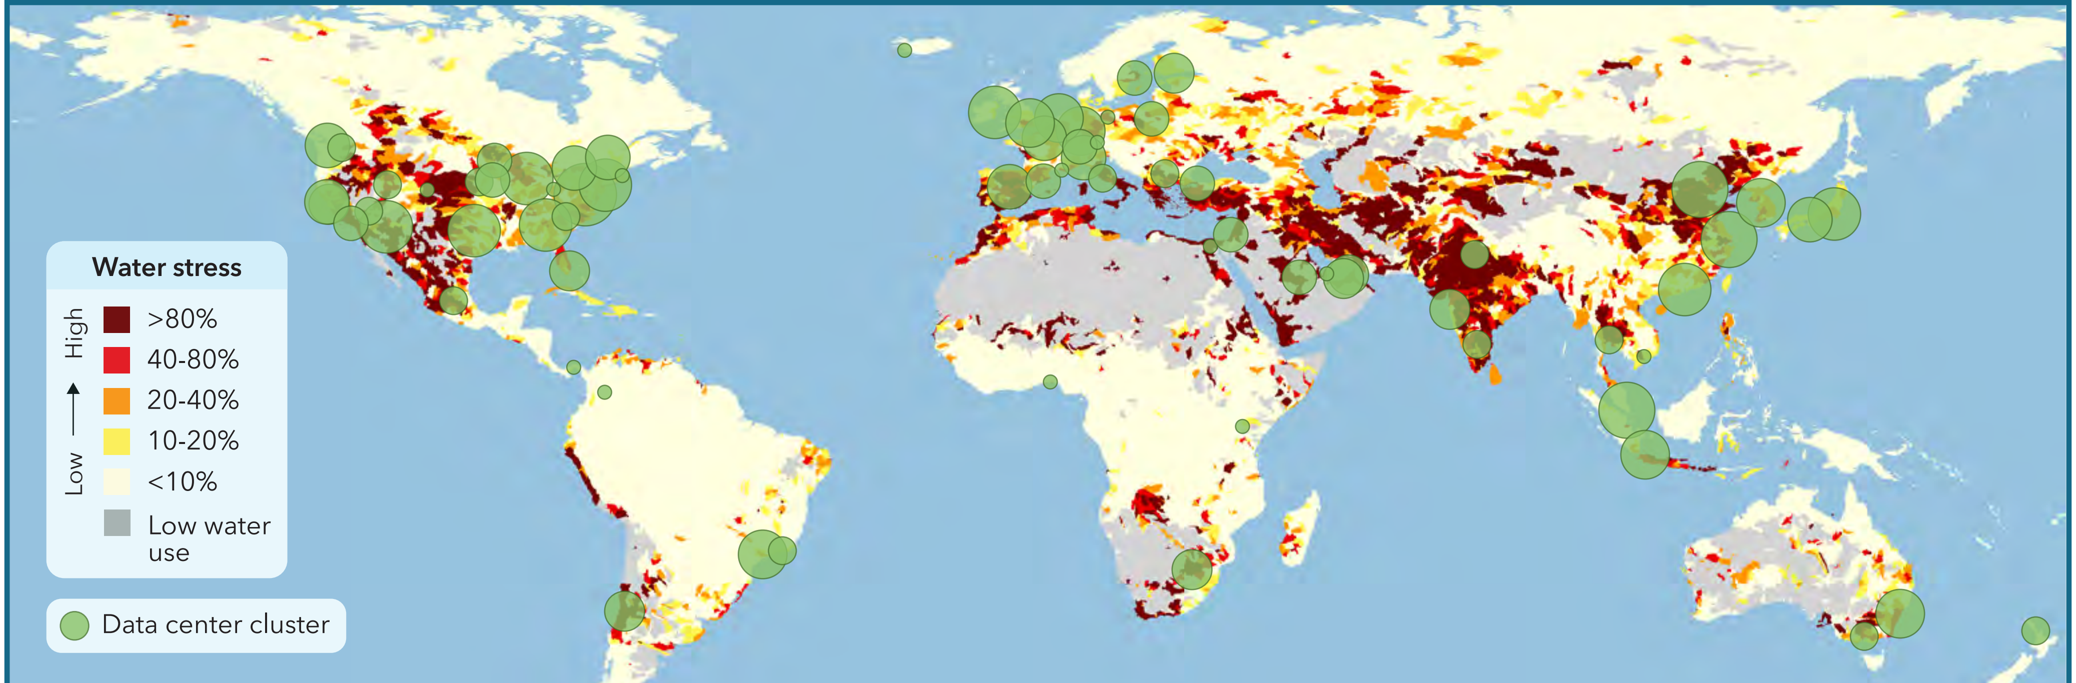


### The changing industrial geography of data centers

Historically, data centers were located near major population centers to minimise latency – the delay between a user's action and a server's response. Proximity to users and the internet backbone was essential for real-time services like search, streaming, and e-commerce. AI data centers, however, have different needs: they focus on large-scale model training and inference, so latency matters less. Site selection is now primarily driven by low energy costs. Cheap real estate and low-cost cooling are also considerations, but the latter is not a major driver in siting. It means that the places that are most likely to see hyperscale data center investment in the future include more rural and water-stressed areas. Current hotspots of investment include Texas, India, and Johor in Malaysia, where growing demands have led to data center specific water tariffs and multiple reclamation projects.

# The water – data – energy nexus

## The growth of renewables is moderating the upstream water footprint of the data center sector

While thermoelectric power generation drives significant indirect water use, the sector's shift to renewables, which require little or no water, could reduce this footprint even as energy demand grows.

Indirect water use occurs at the thermoelectric facilities that power many of the world's data centers. Data centers are increasingly committed to renewable energy, as well as substituting water-intensive coal sources for natural gas. These shifts are likely to lower the overall water intensity of power generation for data centers.

### Sources of power and indirect water use

Power plants use water primarily for cooling, but also to produce steam. The water use associated with the power generation for data centers is counted as indirect water use in total data center water demand.

Within thermoelectric generation, coal remains the thirstiest source, followed by nuclear and then gas, where combined-cycle plants offer significant efficiency gains. Some coal capacity has been extended to meet rising demand, but the oldest, once-through-cooled units are likely to be phased out, with net coal usage for data centers expected to fall by 2030. Gas, meanwhile, is expected to gain share in both grid and data center supply, with net gas usage for data centers expected to rise 127% by 2030. Nuclear power output for data centers is likely to remain broadly stable until small modular reactors begin deployment in the 2030s.

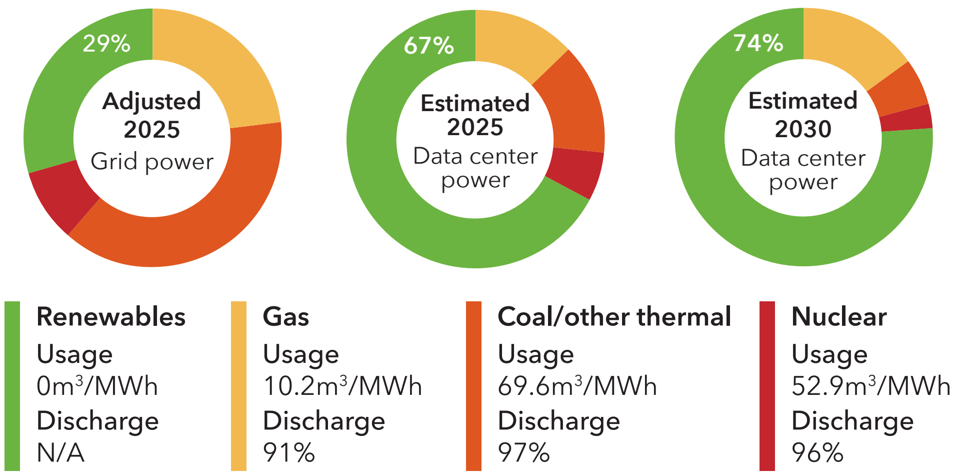


### The data center power mix: accounting for renewables

Most major data center operators – including Amazon, Alphabet, and Microsoft – have now made clear, time-bound commitments to power their data center fleets with renewable energy. However, operators that have yet to reach these goals continue to depend heavily on non-renewable-matched grid electricity, where thermoelectric generation remains dominant and highly water intensive. Because much of the renewable capacity available on the grid has been contracted to data centers through long-term power purchase agreements rather than owned outright, we have adjusted our estimates of the grid mix to exclude data center demand and reflect the higher water content of the remaining supply.

We have accounted for current and expected investment by data center companies in renewable capacity by using an estimated data center-specific power mix to calculate their indirect water intensity. Relying only on regionally sourced grid mixes would ignore their impact on the renewable energy landscape. Data center companies are building onsite capacity and stimulating external investment through power purchase agreements, often paying premiums for renewable grid offsets. While interconnected grid systems make it difficult to isolate the exact source of electricity delivered to any customer, our estimated power mix reflects the sector's positive impact on renewable energy supply.

If renewable capacity for data centers expands more slowly than has been forecasted here, total water use in power generation for data centers will rise faster than expected by 2030.

Power generation for semiconductor fabs is more reliant on thermoelectric sources: while the semiconductor industry accounts for only 15% of total new economy power use, their demand accounts for 22% of total indirect water use in power generation.

# The water – data – energy nexus

## The water impact of the changing power mix

The shift to renewables and gas in the data center power mix will outweigh the near-doubling of energy demand for data centers by 2030, with indirect water consumption through power generation rising only 18%. However, water risks persist since power plants are increasingly reliant on freshwater for cooling.

Onsite cooling and offsite power generation make up total data center water use. As the sector pivots to less water-intensive power sources, this balance will shift. Fuel mix changes are also accompanied by the phasing out of once-through cooling, leading to higher relative consumption and greater use of freshwater.

### Water consumed in data center power supply

In 2025, power supply to data centers used more water than onsite cooling, but as the sector reduces its reliance on thermoelectric power this will change. By 2030, power supply will account for 48% of total data center water consumption.

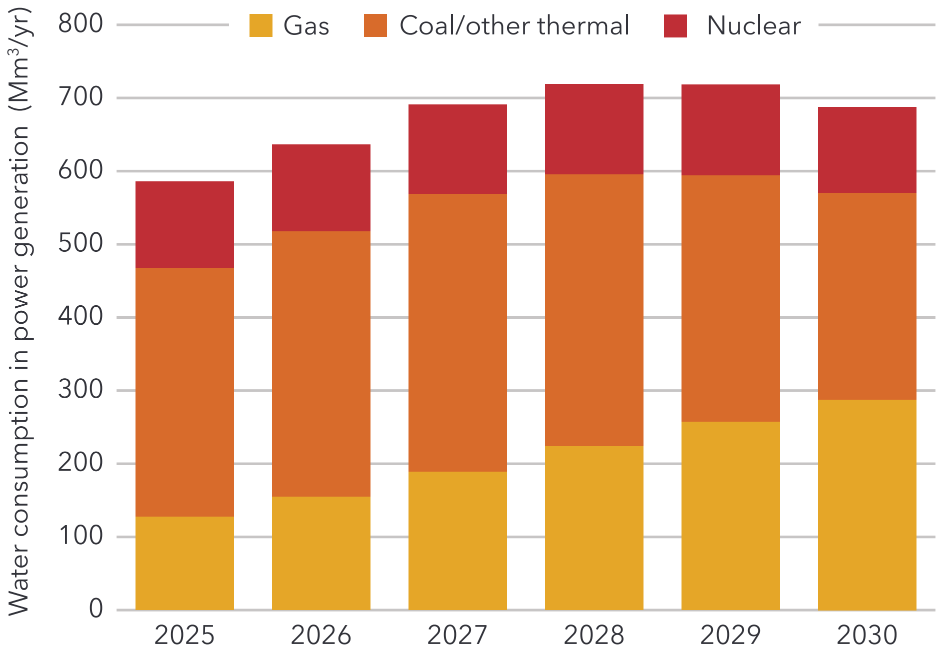


**Power generation as % of total water consumption**

| 2025 | 2026 | 2027 | 2028 | 2029 | 2030 |
|---|---|---|---|---|---|
| **54%** | **53%** | **52%** | **51%** | **49%** | **48%** |

### The data center power mix: accounting for renewables

Thermoelectric power plants are typically cooled in one of two ways:

- **Once-through cooling:** Water (typically seawater or river water) is drawn through heat exchangers to cool the power plant. It is then discharged at a higher temperature to the source water body. This results in significant water withdrawals and discharge, but relatively minimal consumption.
- **Recirculating cooling:** Cooling towers are used to dissipate heat through evaporation. Withdrawals are much lower, but relative consumption is considerably higher. These systems can be adapted for seawater with specific design and treatment alterations, but freshwater is typically used.

### Water risks and strategies for power generation

In recent decades, power plants worldwide have been pivoting to recirculating cooling systems due to the impact on marine life associated with once-through cooling. While withdrawals have fallen, total consumption has remained stable. This also means that power generation is increasingly reliant on freshwater, and therefore more exposed to water stress.

This risk is heightened in regions that depend on water-intensive fuels like coal, such as India and China. For example, a lack of freshwater for cooling cut potential growth in India's power sector by 20% in 2016.

Some newer power plants have installed dry cooling systems to reduce water use. At older plants, however, high costs often prohibit retrofits. Because power plants typically draw from surface water sources, they have limited ability to offset withdrawals through utility network efficiency. There is, however, considerable scope for power plants to use treated wastewater for cooling. This requires system adjustments and long-term infrastructure planning, but can offset future freshwater demands. For example, the largest nuclear plant in the US, Palo Verde in Arizona, has used recycled water for cooling since 1985.

# Closing the supply-demand gap

## Operators need strategies to guarantee their access to water

Semiconductor fabs and data centers face critical water risks, requiring specific strategies to secure sustainable water supplies and mitigate operational risks.

### Water risks in the new economy

Data centers and semiconductor fabs are the heart of the digital economy. But without reliable access to water, chips cannot be produced or cooled, systems fail, and downtime costs soar. Yet water availability rarely tops the list when sites are chosen. Instead data centers and semiconductor fabs develop strategies to ensure that they have the water when they need it. There are three types of risk they face: physical (an extended drought might curtail suppliers), systemic (a utility water system might fail as a result of underinvestment), and political (a community might deny an operator the right to withdraw water).

Increasingly, the preferred means of reducing water risk in both sectors, besides reducing water requirements, is collaboration with local utilities to develop the capacity and resilience of their supplies. This comes as heightened public awareness of each sector's water impact is risking their licenses to operate and threatening project feasibility.

### Semiconductor risk mitigation strategies

#### Physical

1. Optimize processes to minimize ultrapure water usage per wafer and mitigate rising intensity.
2. Partner with wastewater utilities to develop supply systems that can deliver the constant supply needed for fab operations.
3. Develop onsite water reclamation facilities to reduce freshwater demand.

#### Systemic

1. Build resilience through regional water sharing agreements with science parks and industrial clusters.
2. Consider water availability alongside other critical factors during fab site selection.

#### Political

1. Leverage government funding in semiconductor facilities (such as the US and EU chips acts) to invest in sustainable water infrastructure.
2. Explore industrial to municipal or agricultural reuse to improve the license to operate of new fabs.

### Data center risk mitigation strategies

#### Physical

1. Reduce demand by relying on ambient air flow or dry cooling systems.
2. Protect supplies through ownership of senior water rights (in the Western United States).
3. Co-develop water reuse projects with local wastewater utilities to ensure an uncontested supply that does not depend on natural freshwater availability.

#### Systemic

1. Find utility partners who have diversified water resource portfolios.
2. Work with utilities to address operational weaknesses such as leakage and aging critical infrastructure.

#### Political

Commit to water stewardship goals such as Water Positive that promise to replenish within the basin at least as much water as is withdrawn. This may involve financing catchment initiatives such as tree planting or WASH (water, sanitation, and hygiene) projects which do not directly deliver additional water but do ensure the support of water-related NGOs.

# Case studies in resilience

Partnerships are key in securing new supply

Water risk is on the rise in the new economy. Two additional partnership strategies can help close the supply-demand gap: reuse and infrastructure renewal. Here are two examples of these strategies at play.

## Reuse

CASE STUDY

Partnerships can accelerate water reuse, supporting water security.

### Intel in Arizona

#### Context

Intel sought to reduce the impact of its growing water footprint, especially in water-scarce Arizona. The company aims to reach net water positive by 2030.

#### What Intel is doing:

- At its main Arizona campus, Intel entered into a public-private partnership with the City of Chandler to build the Ocotillo Brine Reduction Facility, achieving ~96% water recovery through advanced treatment and brine concentration.
- The utility owns and operates the facility, while Intel funds capital and operating costs, allowing increased reuse capacity without raising municipal water tariffs.

#### Impact

This partnership, along with efforts to lower dissolved solids in reclaimed water, enables the City of Chandler to use treated wastewater for deliveries to the neighbouring Gila River Indian Community. In 2023, onsite reuse led to water savings of 12.8 million m³, and restoration projects added 4.2 million m³ of additional water supply to local communities, reducing the site's impact on municipal supply and offsetting production losses.

## Infrastructure renewal

CASE STUDY

Partnerships can unlock new water supply through network efficiency.

### Amazon in Mexico

#### Context

Mexico City and Monterrey face chronic water stress, with aging infrastructure leading to significant leakage and supply losses. To support its operations in the region, Amazon partnered with local utilities to accelerate smart-water upgrades.

#### What Amazon is doing:

- Amazon is funding the deployment of Xylem's digital water-management technologies, including pressure-monitoring tools, real-time leak detection software, and cloud-based analytics.
- Working with local utilities, the program is modernising distribution networks, speeding up repairs, and reducing non-revenue water.

#### Impact

The initiative is expected to save more than 1.3 million m³ of water annually across the two cities. More of the water produced by utilities will reach residents, reducing strain on municipal systems and enhancing service reliability. Through digitalization and targeted infrastructure upgrades, Amazon is helping improve long-term water security in the region, supporting both community resilience and its operations in the face of rising water stress. Amazon has also decided not to use any water in its data center cooling operations because of the high water stress in the region.

# Towards net zero water

The strategies for building resilience are available and feasible, they just need to be scaled

A CALL TO ACTION

Achieving the 'Water Transition' requires changes in both water supply and demand so that future growth can be decoupled from resource availability.

Climate uncertainty and rising demand mean that improved water resilience will be critical to the development of the new economy. While onsite efficiency gains are included in our water use assumptions, there is still significant scope for new withdrawals to be met with low-impact sources and water intensity to be minimized.

## The three levers of water resilience in the new economy

1. Scale reuse partnerships and smart infrastructure: invest in mutual municipal to industrial reuse and reduce non-revenue water and leakages.
2. Boost onsite efficiency: apply best practices from leading data centers and chip fabs.
3. Optimize the energy mix: support renewable integration to reduce overall water use intensity in power generation.

## Onsite efficiency

Our assumptions for new economy withdrawals by 2050 already account for expected efficiency gains in each sector:

- Cooling system changes in data centers will save 2.4km³ annually by 2050.
- Reuse in chip fabs will save 12.5km³ by 2050.
- The expected growth in renewables for data center and chip fab power supply will save 103.5km³ by 2050.

Investment in utility resilience is therefore the most cost-effective way to neutralize future withdrawal growth. But maintaining these trends in onsite efficiency will still be key in ensuring the success of the Water Transition by preventing unmanageable demand growth.

## The role of each industry in the Water Transition

**Power:** Significant savings are expected to come from the renewable transition. Net zero water can however be accelerated by faster adoption of renewable energy, particularly in regions more heavily dependent on coal.

**Data centers:** The sector is expected to see strong efficiency gains with the adoption of new cooling systems. They can further mitigate their impact by accelerating the adoption of non-potable water sources.

**Semiconductors:** The semiconductor sector will see significant withdrawal increases to 2050. Increasing recycling rates will be counteracted by rising water use intensity as the sector pivots to advanced node production.

Most of the water savings in the Water Transition scenario come from semiconductor fabs. They are expected to have the largest increase in water use by 2050, creating greater potential for additional withdrawals to be offset with the use of recycled water.

## Top-level variables and caveats

There are many uncertainties surrounding the long-term future of the new economy which could impact the trajectory and magnitude of its water impact. They are:

- The rate and intensity of AI adoption in our everyday lives.
- Chip efficiency breakthroughs such as quantum computing and new types of AI-focused chips.
- Compute efficiency gains in AI models meaning that energy intensity is drastically reduced.
- Geographic and temporal variability in the adoption of renewable energy.

Nonetheless, a Water Transition is still necessary to counter the water risks posed by the new economy: proactive action ensures resilience under multiple scenarios.

# The opportunity in shared resilience

## The new economy need not require any additional freshwater withdrawals. The opportunity is in wastewater reuse and in improved network efficiency

Data center operators and semiconductor manufacturers can potentially meet all of their future water needs without requiring any additional freshwater withdrawals – if they are prepared to partner with utilities. Two specific opportunities stand out: the amount of wastewater produced by households that could potentially be treated and made available for reuse, and the amount of freshwater withdrawn and treated by public utilities but subsequently lost as a result of leakages. Together these resources are more than capable of meeting the additional water needs of the AI revolution.

### Wastewater reuse

The biggest opportunity to increase water supply without raising freshwater withdrawals is wastewater reuse. Globally, around 320 km³ of wastewater is produced each year. The cost of reclaiming this depends on the additional treatment and distribution infrastructure required. GWI's analysis suggests up to 160 km³ per year could be supplied at a levelized cost below $1/m³.
As our case studies show, major data center operators and semiconductor manufacturers are willing to invest in wastewater-reuse projects with municipal utilities to secure supply. These schemes also improve overall utility resilience by diversifying sources. Where data center or semiconductor demand is not conveniently located, there is growing interest in supporting "offset" reuse projects that benefit other customers while strengthening system resilience.

### Reclaimable network losses

The second-largest resource available to data center operators and semiconductor manufacturers is improved network efficiency. Globally, about 21% of municipal water withdrawals are lost through leakage and other inefficiencies, creating a potential resource of up to 100 km³. Utilities can access this through better pressure management, leak repair and improved network planning. Some data center operators, such as Microsoft, are already investing in utility efficiency to secure their water supplies. It is a partnership strategy that benefits the data centers as much as the communities that live alongside them.

Increasingly, the technologies used in identifying and addressing network losses are based on AI. In that sense, it is a double partnership.

**The cost of additional water through utility partnerships**

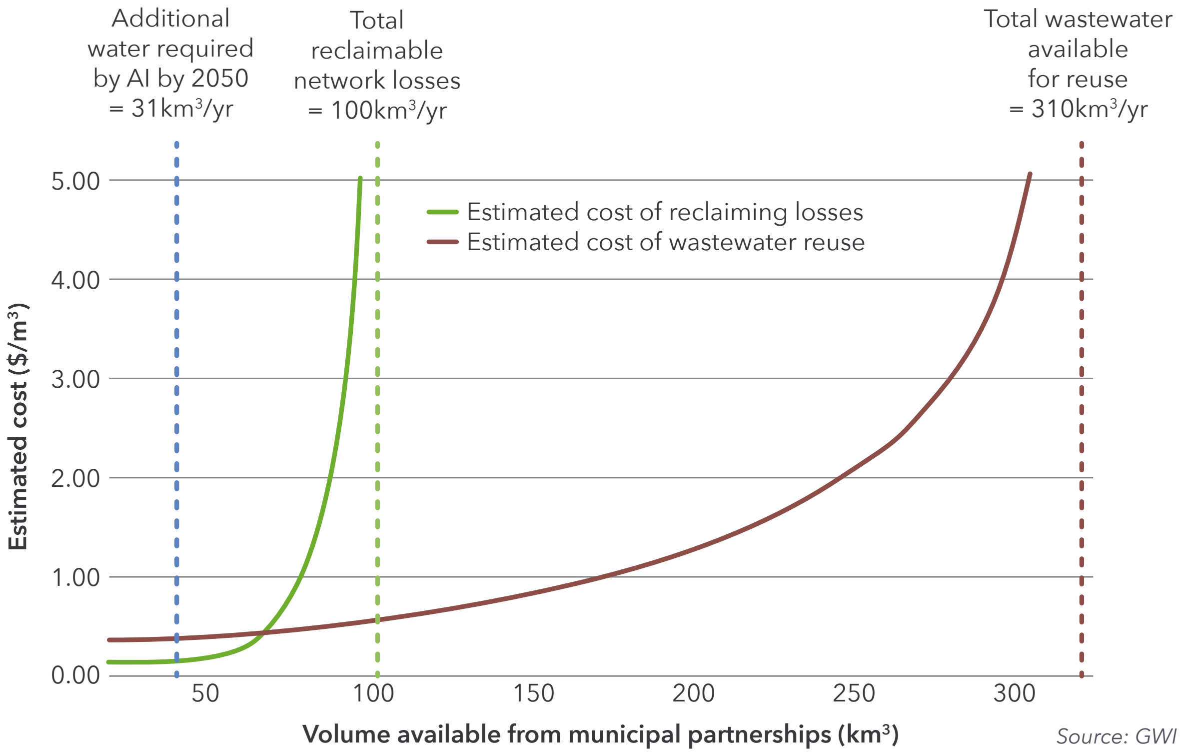


*Source: GWI*

The marginal cost of recovering water from wastewater and network losses increases as the lowest-cost opportunities are exhausted, and capturing the full resource eventually becomes prohibitively expensive. Although the new economy's additional water demand by 2050 is a relatively modest 31 km³, its geographic distribution means some of that supply could cost more than $1/m³.

# Conclusion: a new partnership

Water enables the new economy and the new economy can enable the environmental transition

The new economy offers a powerful starting point for collaboration between industry and local communities on water management. It faces significant public scrutiny over its water use, and the data center and semiconductor industries have already made some of the strongest commitments to reducing withdrawals and consumption. These pressures are driving innovation and partnership. In time, the models now emerging around data centers and fabs can form a blueprint for wider cooperation between cities and industries to navigate the new realities of a changing water cycle.

## What the new economy needs

**Growth:** It cannot allow water to hold back growth. Our future economic wellbeing depends on the productivity gains of the new economy.

**Reliability:** The new economy is a highly controlled 24/7 phenomenon. It cannot allow the new extremes of the water cycle to disrupt its operations.

**License to operate:** Ultimately, communities have the power to close businesses that don't respect what they value.

## What cities need

**Reliability:** Many urban water systems are unreliable after years of underinvestment and struggle to meet new water supply demands.

**Resilience:** The growing intensity of the extremes of climate change means that cities around the world are ever more vulnerable to floods and droughts.

**Harmony:** A successful city needs to balance the needs of its population and the needs of the businesses that provide jobs and prosperity.

## What the new economy has to offer

**Investment:** Water security has a clear economic value to businesses in the new economy. It means that they are usually prepared to invest in public infrastructure to guarantee dependability of supply.

**Intelligence:** Artificial intelligence can potentially revolutionize the way in which we manage water. It will introduce an aspect of agility to the way in which we respond to the changing water cycle.

**Commitments:** CSR commitments to zero net water consumption create the entry point for municipalities to collaborate with businesses to develop shared water assets.

## What cities have to offer

**Infrastructure:** Cities are responsible for delivering shared water and wastewater infrastructure. They can offer the broadest menu of water resource solutions to businesses located in their region.

**Opportunity:** Businesses looking to reduce freshwater withdrawal while increasing process water use need to see the opportunity in municipal water reuse, leakage reduction, and improved efficiency and reliability.

**Consent:** In the long term, aligning the water interests of new economy businesses with the communities in which they are located is the best guarantee of economic, social and environmental sustainability.

## How they can work together

The new economy and cities share both risk and opportunity. By aligning investment in intelligent infrastructure and unconventional water resource development, they can build integrated water systems that protect communities and secure growth. Collaboration on reuse, data, and resilience can transform water from a constraint into an enabler of sustainable innovation, turning local water stewardship into a shared foundation for global progress.

# Appendices

The following appendices provide supporting data, detailed methodologies, and additional context to complement the analysis presented in this whitepaper.

**Overall sector scope and data methodology**

**Data centers:** Includes all global data centers (hyperscale, colocation, enterprise). Excludes bitcoin mining. Water use figures are based on net utilized MW estimates and forecasts, combined with GWI's forecast of regional cooling system breakdowns. Power capacity and use data comes from 451 Research (S&P Global), *Datacenter Services & Infrastructure Market Monitor & Forecast*. Climatic and seasonal factors are modeled and combined with the above to generate regional and sub-regional WUE estimates which are then combined with regional and sub-regional utilized data center capacity data to generate water use estimates.

**Semiconductors:** Includes global front-end (wafer and semiconductor) manufacturing. This includes production of semiconductors for data centers, consumer electronics, automotive and other industrial applications. Non-integrated circuit applications (such as solar photovoltaic manufacturing), as well as upstream silicon and other raw material production, is excluded. Water use data is based on company reporting and internal estimates. Market growth estimates are based on assumptions around data center expansion and the growing share of data centers as a proportion of semiconductor output and market revenue.

**Power generation:** Includes global power generation for data centers and semiconductor fabs. Data center power mix is based on estimates derived from company reporting (see page 15 methodology for further information). Semiconductor power mix is based on estimates derived from company reporting, weighted to the market share of each company. Energy intensity increases as a result of the move to more advanced nodes. See average intensity per semiconductor layer (page 8 methodology) for this breakdown. Manufacturing energy intensity by node size is based on IEA (2025), *Energy and AI*.

*Figure 1. New economy water withdrawals by sector, km³*

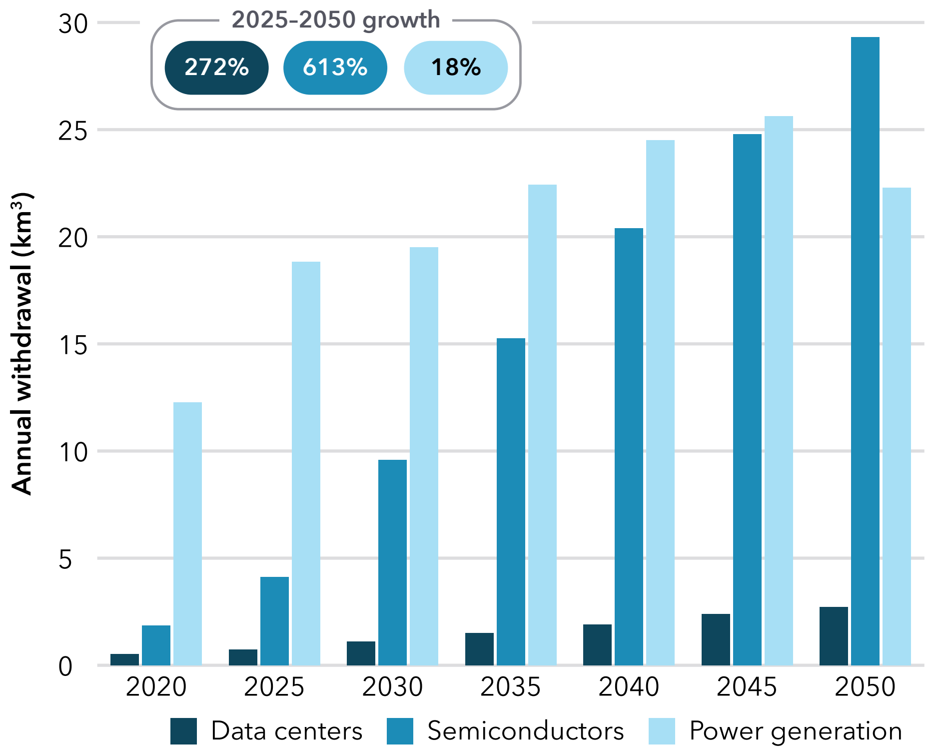


These 3 sectors form the new economy. Within this, AI is a focus since it is the most significant driver of growth in all 3. The new economy can be considered as the AI value chain. The shift to renewables, as well as broader digitization trends are also considered as drivers of the new economy. All water stress metrics and maps use the World Resource Institute's *Aqueduct* (Water Risk Atlas) baseline (current) water stress data, at an annual temporal resolution.

**Page 4: How can we limit future withdrawals to today's levels.**

This chart illustrates how reclaimed network losses and the reuse of municipal wastewater can offset additional withdrawals driven by growth in the new economy. Savings from reclaimed losses and improved network efficiency are estimated based on each sector's reliance on municipal versus self-supply and average network loss rates in regions where these sectors are most prevalent. The potential volume of reused wastewater is calculated from the extent of internal reuse possible and the substitution opportunities for offsite reuse and unconventional supply.

# Appendices

**Page 6: The water impact of 30 minutes of AI use**
What is counted as AI use? Our 32.5Wh per 30 minute figure is based on intensive AI usage (such as generating code or deep research), rather than casual conversations. It is counted as 15 'time-scaled' queries per hour (7.5 queries/30 minutes), using the power consumption of AI inference assumptions from Microsoft (2025), *Energy Use of AI Inference*, with each query consuming 4.3Wh of electricity. 'Time-scaled' queries weight AI inference requests by how long compute resources are used rather than only the number of queries, accounting for factors such as reasoning complexity. Using OpenAI's estimate that one ChatGPT query uses 0.34Wh, our assumption of 32.5Wh corresponds to 96 'normal' queries over 30 minutes, equating to intensive usage.
To estimate the water use embedded in chip manufacturing scaled down to 30 minutes of AI usage, we used assumptions around the required chip utilization to process these queries and internal estimates on the water intensity of advanced semiconductor production. Chip lifetime was assumed to be 2.5 years, allowing amortisation down to 30 minutes. These numbers do not include the water embedded in personal devices (e.g. laptops or mobiles) needed to connect to an AI service.
Offsite power generation and data center cooling figures were calculated based on the energy consumption data above, combined with average water use efficiency for power generation supplying data centers (see page 15 methodology), and average water use effectiveness for data center onsite cooling (see page 11 methodology).

**Page 8: Average intensity per semiconductor layer**
This graph indicates that as the market moves to more advanced nodes to fulfill the compute demands of AI, total water intensity for the semiconductor sector will increase. Water intensity by node size is based on M. Garcia Bardon et. al. (2020), *Cradle-to-gate Life Cycle Assessment of CMOS Logic Technologies*. GWI estimates on the changing market breakdown by node size can be seen in figure 2. The intensity factors and node size breakdown are combined to calculate the total relative intensity of semiconductor manufacturing over time, expressed as $m^3$ of water consumption per 300mm wafer layer. Total intensity increases with smaller nodes due to the need for more wafer layers in denser, more advanced chips. 10nm nodes were used as the threshold between advanced and non-advanced chips. GPUs and other AI-specialized chips typically use a node size of 5nm and below. 7nm and 10nm chips will still support AI growth in cloud and lesser advanced AI applications.
Average intensity per semiconductor layer does not include indirect use in power generation. Water use factors in semiconductor yield and is calculated on the basis that a majority of spent ultrapure water is reused in cooling towers.

*Figure 2: Node size market breakdown for the global semiconductor market*

| | 2020 | 2025 | 2030 | 2035 | 2040 | 2045 | 2050 |
|---|---|---|---|---|---|---|---|
| >10nm (all non-advanced) | 91.0% | 87.5% | 76.5% | 67.0% | 58.0% | 48.5% | 39.0% |
| 10nm | 3.0% | 2.0% | 3.5% | 5.5% | 7.0% | 8.5% | 10.0% |
| 7nm | 2.5% | 4.0% | 7.0% | 10.0% | 12.5% | 16.3% | 20.0% |
| 5nm | 2.5% | 4.0% | 7.0% | 10.0% | 13.5% | 16.3% | 19.0% |
| ≤3nm | 1.0% | 2.5% | 6.0% | 7.5% | 9.0% | 10.5% | 12.0% |

**Page 8: Semiconductor fab locations**
- This data focuses on fabs producing advanced chips (mostly 10nm and below for those where node size is published) and foundries producing advanced (300mm) wafers.
- For North America, existing fab locations come from the Semiconductor Industry Association. Planned fab data comes from tracked project announcements.
- For East Asia, existing and planned fab location data comes from tracked project announcements.

**Page 10: The water use effectiveness of cooling technologies**
Average WUE for cooling tower and adiabatic systems is based on Lawrence Berkeley National Laboratory (2024), *United States Data Center Energy Usage Report*. WUE for direct-to-chip systems is based on our estimates on the breakdown of facility level cooling systems currently deployed with direct-to-chip systems. The split between discharge and consumption is based on the evaporation rates of different cooling technologies in varying climates. Average WUE varies from real-world or applied WUE since different cooling systems are impacted by climatic and seasonal factors in different ways. For instance, a cooling tower typically runs much more often than an adiabatic system, which will operate for only 2 months of the year at a maximum. Humidity and temperature are also important. Humidity reduces the difference between ambient dry-bulb and wet-bulb temperatures, meaning that the surrounding air can absorb less heat, reducing cooling efficiency and leading to in-

creased water use. Consequently, cooling tower systems in hot and humid areas have higher WUEs. Adiabatic systems, meanwhile, are less effective in humid regions and are typically avoided. These dynamics are accounted for in our cooling system breakdown and forecast of water use in the sector.

*Figure 3: North America data center water withdrawal seasonality with cooling system breakdown*

| | Jan | Feb | Mar | Apr | May | Jun | Jul | Aug | Sep | Oct | Nov | Dec |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Tower | 3.56 | 5.33 | 10.67 | 20.14 | 36.58 | 70.71 | 102.96 | 79.13 | 24.25 | 16.29 | 10.67 | 4.44 |
| Adiabatic | - | - | - | 0.58 | 1.02 | 6.7 | 9.62 | 8.01 | 2.91 | 0.29 | - | - |

**Page 10: The seasonality of water use in North American data centers**
This chart estimates monthly water use in North American data centers, excluding Mexico, based on cooling tower and adiabatic system performance across different climates. Monthly water use was modeled using average temperatures for major climate zones, accounting for how each technology responds to dry-bulb and wet-bulb conditions (the difference between ambient air temperature, and the lowest temperature air can reach via evaporative cooling). Combining these factors with the regional cooling system mix produces a month-by-month water use estimate, seen in figure 3. Adiabatic systems operate only from April to September, with one-third of their total withdrawals occurring in July.

**Page 11: Cooling system trends, 2025–30**
This breakdown of cooling systems in the global data center market comes from GWI's forecast of water use in the sector. The distinction between facility and chip-level cooling systems is made by the green and pink areas. Chip-level systems (conventional air handlers or direct liquid) absorb heat from the servers themselves while facility systems (adiabatic systems, cooling towers, air chillers) reject this heat out of the data center to the environment.
Cooling system breakdowns at a regional, sub-regional (e.g. different parts of the USA) and country-level for 2025 and 2030 were estimated based on available market data, interviews with industry experts, and assumptions around the use of different cooling systems in different climates and regions. The global breakdown is based on the market share (by installed MW capacity) of different regions, sub-regions and countries, and their respective cooling system breakdown.

**Page 11: Forecast data center power and water use 2025–2050**
The forecast of power use is based on estimates of power use up to 2050, assuming widespread adoption of AI, and a continued buildout of data centers.
The water use data is based on GWI's data center water use forecast up to 2030, and thereafter on estimates of WUE across the data center market. It is expected that the period of most intensive change in cooling systems (and declining WUE) will be between 2025 and 2030, as seen in figure 4.
Global WUE is expected to plateau around 2040–50 (figure 4). This is for a variety of reasons:

- Direct-to-chip cooling will reduce water demand in the long term (due to higher allowable chip temperatures and greater ease of cooling with air at facility level), but not all data center workloads will require direct liquid cooling at the chip level, and will continue to use conventional systems for cost reasons.
- Cooling towers and other evaporative cooling systems will still be used in areas of high humidity and water availability due to greater energy efficiency, especially as PUE (power use effectiveness) is increasingly regulated in the sector.

*Figure 4: Global data center WUE (water use effectiveness), measured in water withdrawal per KWh of data center compute.*

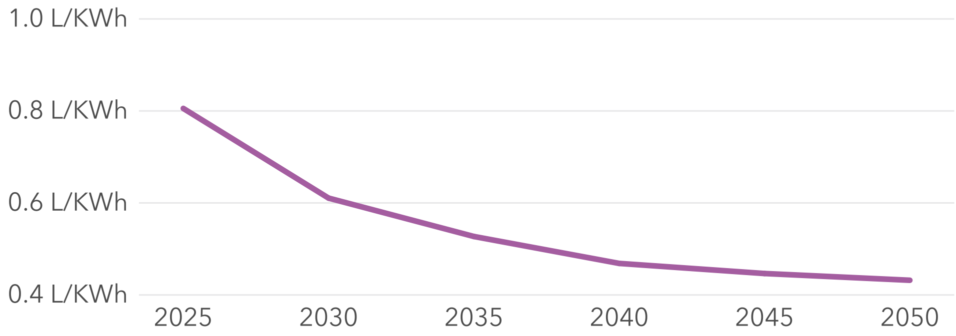

# Appendices

**Page 12: Water use in daily life**

**Cup of Coffee:**

- Direct usage: the water in the final drink.
- Power: the water withdrawn in generating the energy needed to process the coffee beans and produce hot water.
- Embedded: the agricultural water use (majority green water), based on the assumption that a typical cup of coffee needs 7g of roasted beans (from 11g of green beans), and that to produce 1kg of green beans, around 17,000 liters of water is needed.

**Phone browsing:**

- Power: the power use figures represent the water withdrawn to generate the power to charge the phone, as well as to run the cloud data centers that support internet browsing and mobile apps.
- Cooling: the water withdrawn onsite at data centers to cool the cloud data centers that support internet browsing and mobile apps. This is derived from the estimated data center power mix.
- Device-embedded: this figure (2.2L) is derived from an estimated 12,000L embodied water footprint for a typical mobile phone. This was allocated over an assumed 3-year lifetime and 3 hours of daily use, giving a per-hour share of manufacturing water that scales to approximately 2.2L for a 30-minute session.

**General activity:**

- Direct use: drinking water and other direct uses.
- Power: water used indirectly through electricity generation for general activity such as lights and appliances, based on the adjusted grid mix.
- Embedded: accounts for the indirect and embodied water use of goods and services, including food, clothing, consumer products, scaled from their respective lifecycles to 30 minutes.

**AI use:**

- Power: this figure represents the water withdrawal required to generate the power needed to fuel 30 minutes of AI use. The cooling figure represents the water withdrawn onsite at data centers processing AI requests and running model inference for cooling purposes. The power and cooling figures here come from the 30 minutes of AI use figures on p.5.
- Desktop computer: this figure (7L per 30-minute session) is based on an adjusted embodied water footprint of ~7,500L. This adjustment accounted for the agreed weight of a modern desktop relative to older reference models. The footprint was then allocated across an assumed 5-year lifetime and 4 hours of daily use, with a small uplift applied to reflect additional components. The inclusion of the water footprint of a desktop computer explains the higher overall footprint for 30 minutes of AI use here than on page 6, where the water footprint is measured from the isolated perspective of the data center, and excludes the personal device.
- Embedded in device: this represents embodied manufacturing water, apportioned over typical usage patterns, and are intended as indicative comparative values rather than precise engineering measures.
- Out of basin refers to the proportion of water use that occurs outside of the basin in which the final product is used or manufactured.

**Page 12: Water use in data centers versus other industries**

The water use per dollar of revenue data is calculated through GWI estimates of water consumption in each industry. This was combined with estimates of global revenue for each sector to come up with a unit water use per dollar of revenue generated. The power data here includes the global power generation sector, and is based on unadjusted grid breakdowns, not the adjusted mixes shown on page 15.

**Page 13: The water use effectiveness of the big five**

Water use data comes from company sustainability reports. We have focused on the 'big five' hyperscale data center operators (rather than including co-location providers, which also account for a large portion of total capacity) for a variety of reasons:

- They are at the heart of the data center water use debate and attract the most public scrutiny.
- They are more likely to report water use metrics.
- They are more closely and directly involved in water use strategy than co-location providers, whose cooling and water use strategies are more variable.

While Google does not disclose WUE in their annual sustainability reports, they do disclose a figure of 1.15L/KWh for their AI-related facilities in Google (2025), *Measuring the environmental impact of delivering AI at Google scale*. This figure matches a calculation of their WUE based on disclosed data on water and energy use at all of their data centers.

**Page 14: Where are the biggest data center clusters in the world?**

These data center cluster locations and sizes are based on estimates of the relative size of well-known data center markets. Bubble sizes are illustrative of estimated data center capacity rather than specific MW values.

**Page 15: Water – data – energy nexus: adjusted grid methodology**
The adjusted grid mix separates data center electricity demand from the general grid to account for large-scale renewable procurement by data center operators. Data center energy use by fuel was calculated on the reported power sources for major data center operators (including both hyperscale and co-location). This includes onsite generation of renewable energy (such as adjacent solar farms or onsite wind turbines), renewable power purchase agreements (which helps to finance investment in renewable capacity), and renewable energy certificates, which represent proof that a unit of energy was generated from a renewable source. This data was aggregated for a representative group of major data center operators (who publicly report these values) to estimate a power mix for the global data center sector.
Total data center demand and its associated renewable supply were then removed from the IEA 2024 global grid mix (IEA (2024), World Energy Outlook), and the remaining generation was renormalized to reflect the higher thermoelectric share serving other grid users. This avoids double counting renewables and better represents the indirect water intensity of grid electricity.
The 2030 data center power mix was based on the current rate of renewable capacity build out, and the future commitments of data center operators to expand their use of renewable energy.

**Page 16: Water used in data center power supply**
The chart on page 16 shows the indirect water consumption of power generation to fuel data centers. If this is expanded to account for total withdrawals, then power generation accounts for a much larger proportion of total water use for data centers, as seen in figure 5. This is because power generation water use is much more withdrawal-intensive compared to data centers due to their use of once-through cooling. Power generation as a proportion of total data center water footprint therefore varies significantly based on the metric used.
Semiconductor proportion of total power generation and indirect water use
Semiconductor power use has been included in our estimates of total power generation water use in the new economy. Figure 6 shows the breakdown between data center and fab energy demand and indirect withdrawals through power generation. Data center power use is driven by AI and cloud infrastructure buildout. Similarly, semiconductor power use is pushed up by industry expansion, as well as increasing energy intensity as the sector pivots to advanced nodes. Sector wide energy intensity for semiconductors is expected to increase 90% by 2050, pushing up total energy demand and indirect water use.

*Figure 5: Data center indirect water withdrawals in power generation, by fuel source*

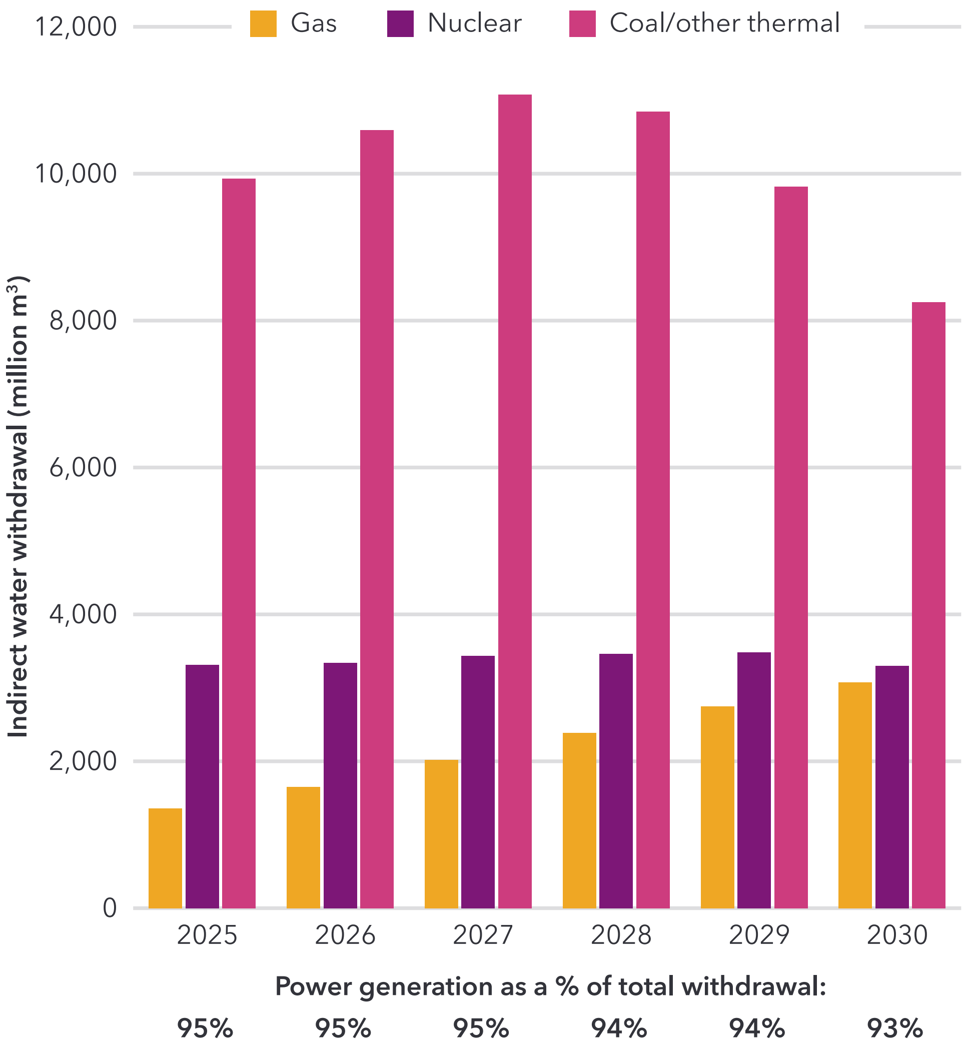

**Page 19: The cost of additional water through utility partnerships**

Total availability is estimated using a projected global municipal freshwater withdrawal of 470 km³ per year (based on Aquastat trends). Around 21% of this (approximately 100 km³) is lost through network leakage. Of the remaining ~370 km³ delivered to households, about 86% returns as wastewater (with consumptive use at roughly 14%), yielding roughly 320 km³ potentially available for reuse.
The marginal cost curve for wastewater recycling is derived from GWI's analysis of capital and operating costs from projects in the DesalData Water Reuse Project Inventory. It assumes treatment to tertiary standards with an additional polishing step, rather than the full triple-barrier typically required for potable reuse. This produces a low-end cost of about $0.40/m³, with costs increasing primarily due to the collection and distribution infrastructure needed to deliver reclaimed water where it is required. These infrastructure-driven costs vary significantly by geography.

The cost curve for reclaiming network losses draws on global cost data from leakage-reduction programmes. At the low end, smart pressure management may be sufficient to reduce losses; at the high end, costs approach those of full pipeline replacement. As physical losses fall below ~4%, further reductions become prohibitively expensive due to the disproportionate cost of locating and eliminating the remaining leaks.

*Figure 6: Data center and fab breakdown between total energy demand (TWh) and total indirect water withdrawal in power generation*

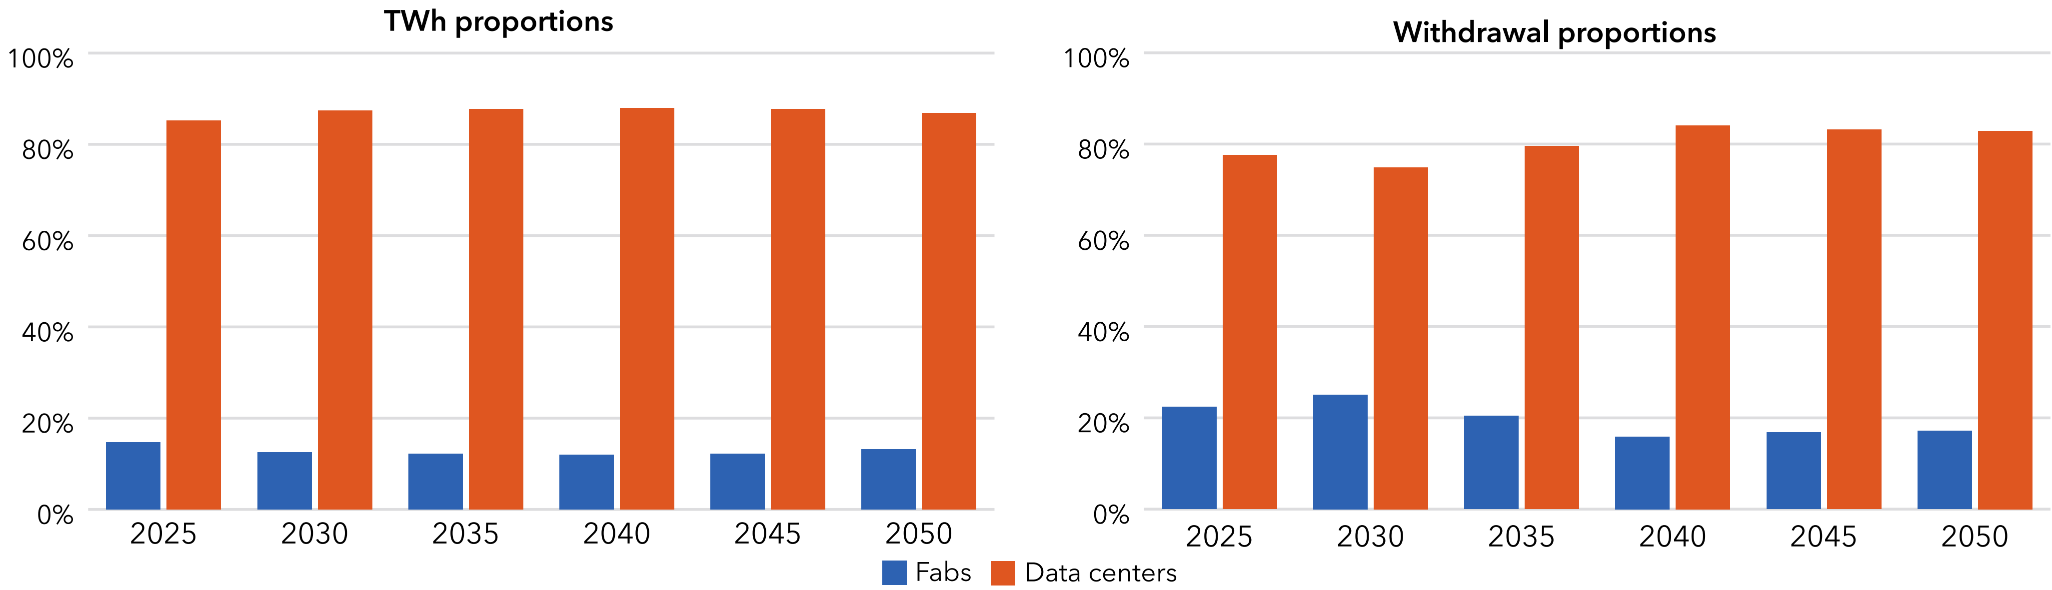

## Watering the new economy
### Managing the impacts of the AI revolution

Written by: Christopher Gasson & Victor Smith
Comment piece by: Matthew Pine
Edited by: Charlotte Oakes
Proofread by: Alexander Truuvert, Ashi Banerjee, Daniil Antonov, Mariana Quaresma, Oluwadara Adebobola, Sam Dennis & Sebastian Hancox
Designed by: Amélie Robert, Christopher Gasson & Victor Smith
Cover design by: Beckstar Design

**Published by:**

Media Analytics Ltd.,
Suite C, Kingsmead House,
Oxpens Road,
Oxford OX1 1XX,
United Kingdom

Tel: +44 1865 204 208



GLOBAL WATER INTELLIGENCE