International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(17s) 2026 pp. 375-384
https://doi.org/10.70917/ijcisim-2026-4746
Article

Synergistic Cooling and Sovereignty: A Socio-Technical Analysis of Deep Ocean Water for AI Data Centers and Desalination in Water-Stressed Developing Regions

Karishma Verma

Independent Researcher, USA

Abstract: Concerns are raised among water-poor developing nations about the growing demand for water by large data centers when evaporative cooling can compete against domestic, agricultural, and ecological water needs. Pumping Deep Ocean Water from thermohaline-stratified layers with average year-round temperatures just above 0 degrees has no need to draw down freshwater resources and is therefore a hydrologically neutral cooling solution. DOW extraction colocated with RO desalination can reduce capital costs for high-value colocated missions and RO energy intensity by using excess waste cold energy extracted from the data center cooling. Risks from discharge can be reduced through accurate thermal and saline plume modeling, multiport diffusers, and adaptive sensor-based discharge management that can simultaneously enable near-coastal marine co-benefits. Resource Offtake Agreements and community governance structures codifying equal allocations as non-negotiable hard law, and a hierarchical multi-agent reinforcement learning optimization layer for conflicting resource demands, are design principles through which community water welfare can be catalyzed by AI-enabled infrastructure for water sovereignty in coastal developing metropolitan areas.
Keywords: Deep Ocean Water Cooling, AI Data Center Sustainability, Seawater Desalination Co-Location, Water Sovereignty Governance, Reinforcement Learning Nexus Optimization

1. Introduction

One new and, to date, relatively unexplored friction associated with AI infrastructure expansion is between digital development and resource security, with governments across the global South competing to attract data center investment as vehicles for broader agendas of economic modernization. The environmental costs (especially water use) of these data centers and how they are in conflict with the domestic, agricultural and ecosystem water needs of local inhabitants are becoming increasingly apparent. This study argues that a technically and institutionally designed alternative is well within our engineering capabilities, and this outcome is not inevitable. Yet the speed of this infrastructure rollout creates a paradox where many water-scarce developing countries are competing with the very large AI systems they want to attract in terms of fresh water as their digital modernization platforms converge .

As with most data centers, the main heat rejection systems consist of evaporative cooling towers and chilled water systems using the evaporation of large amounts of water to offset the heat generated by the data center. The data center industry has grown rapidly, mainly clustered around coastal, semi-arid and developing countries with governments that are eager to stimulate the new digital economy. Water access is a major constraint in parts of the world. Evaporative cooling for data centers thus directly and disproportionately weakens water security for the local, domestic and environmental purposes it serves, and even more so in the Global South, where water access inequality is pronounced [1].

This is a meaningful amount, and data centers are among the most water-intensive built infrastructure typologies for modern economies. Most of this water is for in situ cooling. Unlike electricity demand, water use for in situ cooling is a baseload demand, which occurs regardless of the season, drought or stress on municipal systems. [1]. In water-scarce developing countries, where groundwater and rain are already limited, no cap on water consumption with a large infrastructural investment creates a dual problem of water competition and a disincentive for digital infrastructure investment.

This battle will require radically rethinking where, how and by whom the physical and virtual infrastructure of AI is sited and embedded within local ecosystems of resource provisioning. Plugging a data center into a municipal water grid is just not tenable for coastal communities and economies already overexposed to global warming's climate risks.

Deep Ocean Water (DOW) is below the level of thermohaline stratification, where the temperature is virtually constant (typically 4 to 5 degrees Celsius year-round). Such seawater is attractive, as it forms a huge and virtually inexhaustible thermal sink, requiring negligible water consumption and no energy-intensive refrigeration system to maintain this stable temperature differential. Cold seawater is pumped through a heat exchanger that is connected to the data center's cooling loop, extracting heat from the servers and transferring it to the ocean with a fraction of the energy and water of conventional cooling systems. Since the water is returned to the ocean, this process uses no freshwater. As a result, it has no impact on surface freshwater systems [2].

In addition to cooling, the DOW extraction infrastructure can be used in combination with reverse osmosis desalination plants, as the same pipe that feeds the data center cooling loop can double as a source of pre-cooling seawater for a low-energy desalination plant. With this co-location model, the capital expenditure of the deep-sea infrastructure is amortized over two commercially and socially valuable end-uses, improving the techno-economics of both. The residual cold energy in the seawater following the cooling loop is known to reduce the cost of thermal pretreatment prior to RO. By deploying this circular resource architecture, we can turn the AI hardware from a water sink to a safe drinking water co-producer and ensure that our digital investments are development dividends [2].

In looking at techno-economic, environmental, and governance factors of co-location, environmental discharges, fair water sharing, and nexus optimization through AI, the paper concludes that the AI infrastructure does not have to be a competitor for water in developing contexts. With the right governance and design, it can act as a structural catalyst of water sovereignty, opening up access to the digital economy and serving as another source of clean water supply.

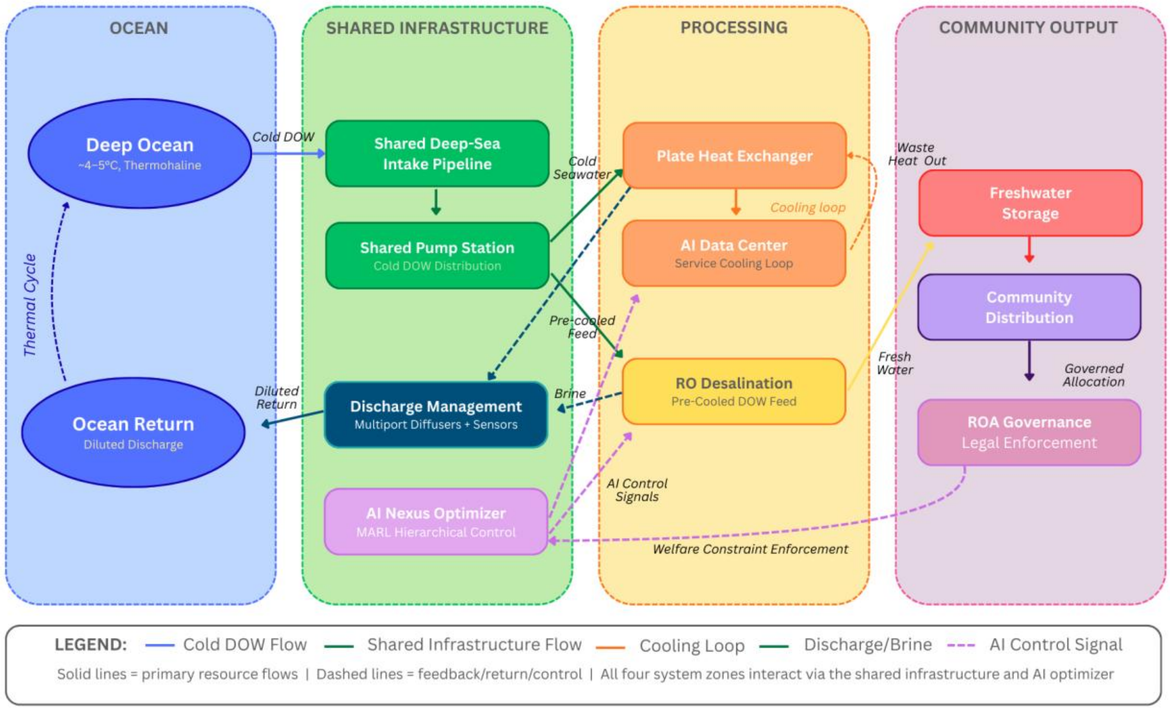

Figure 1: Deep Ocean Water (DOW) Co-Location System: Integrated Resource Flow

Figure 1 provides a systems-level overview of the integrated DOW co-location architecture, illustrating the resource flows and governance linkages discussed in the sections that follow.

2. Techno-Economic Case for Co-Location

There are two economically sustainable ways in which the co-location model under DOW works. One is the sharing of expensive deep-sea infrastructure between cooling of data center and desalination. This is a cost-sharing argument with an immediate impact on project financing. The second is a circular resource economy, where thermal by-products of one system become productive inputs to the other, producing compounded savings over the life of the project. The combination of these mechanisms makes the co-location configuration cheaper than simply deploying the two systems separately and potentially transformative with respect to the financial viability of desalination in the capital-constrained developing country context.

2.1 Shared Infrastructure as a Cost Lever

The most obvious and measurable economic benefit of co-locating AI data centers with DOW-fed desalination plants is the mutualization of two capital-intensive technologies. In any DOW application, be it cooling or desalination, the deep-sea intake pipeline and marine civil works would constitute the majority of the project capital costs. The submarine pipeline, anchoring system, pump station, and heat exchanger manifolds comprise a single system, supporting both end-use applications. The different end-use applications receive cold seawater from the same infrastructure onshore, with the exception of an RO desalination intake where the submarine pipeline has been deployed for a single end use (e.g., the data center cooling loop) from the outset. System CapEx can be amortized over two (instead of one) end-use applications, increasing each end use's IRR [3].

A stand-alone seawater RO desalination plant of any important capacity would need its own seawater intake and outfall and also require meaningful energy-intensive pretreatment to bring raw seawater to RO specification. Pre-treatment typically consists of coarse screens, fine filters, and thermal pre-treatment. Unless pretreated and recirculated on-site, a data center using surface water for cooling in an evaporative system and then evaporating water into the atmosphere would compete with residential, industrial or agricultural uses for this water. The co-location model therefore has an economic and operational rationale, as there is a lower capital outlay in infrastructure, and a single pipeline designed to meet the combined hydraulic needs of both uses has lower unit costs of cold seawater delivery than two distinct pipelines, despite the larger pipeline diameter and associated pumping cost [3].

Besides pipeline sharing, co-location allows sharing of pump house buildings, backup power generation, civil foundations, marine environmental monitoring, and other onshore requirements. This reduces CAPEX for each facility directly and compounds all the way up the infrastructure stack. Certainly in emerging countries where project finance capital markets may be limited and the cost of capital is high, these CAPEX savings can be the difference between financial close and cancellation.

Infrastructure Component	Standalone Data Center	Standalone RO Desalination	Co-Located DOW System
Deep-Sea Intake Pipeline	Not applicable	Full cost (single use)	Shared across both systems
Pump Station	Independent unit	Independent unit	Single shared unit
Heat Exchanger Manifold	Independent unit	Not applicable	Integrated shared unit
Onshore Civil Works	Independent unit	Independent unit	Consolidated single unit
Backup Power System	Independent unit	Independent unit	Single shared unit
Environmental Monitoring	Independent unit	Independent unit	Single shared network

Table 1: Comparative Infrastructure Cost Components: Standalone vs. Co-Located DOW Systems [3]

2.2 Circular Resource Economy

Apart from their economic benefits, co-locations' circularity comes from the thermal and material by-product synergies between the two systems, which lower their operating and environmental costs by closing the material and energy spin cycles between both systems through material recovery and energy utilization.

Thermodynamic circularity describes a process where cold seawater from the deep ocean is pumped through a plate heat exchanger connected to the data center's internal cooling system. As the seawater passes through the heat exchanger, it absorbs heat energy from the computer servers and exits at an elevated temperature. The heat exchanger discharge water is not released into the sea but used as feed water for a co-located RO desalination plant's pretreatment system. Since performance and energy requirements of RO membranes are strongly affected by feed water temperature, using pre-chilled DOW discharge water as feed water for RO units reduces the specific energy requirement for desalination, because less feed water needs to be pumped at low pressure for a given permeate flow [4].

Seawater desalination by RO is among the most energy-demanding water treatment technologies. Energy costs for seawater desalination are often the largest operating costs incurred by an RO plant. Thus, reduction of energy consumption, and in turn of water production costs and water price charged to the end users, represents a direct economic benefit of DOW pre-cooling. This cost saving is of particular interest for a plant that is being constructed to provide water to poorer communities. Deep ocean source water has less biological and particulate material than surface ocean water, which means that membrane fouling occurs less frequently and less severely, thus prolonging the life of the membrane and providing a further operating cost benefit.

The energy requirements of pumping to keep DOW at a depth become a meaningful proportion of the operating costs of both cooling and desalination for both stand-alone and dual-purpose desalination plants. If co-located, some of the pumping energy costs are assigned to the data center's cooling system. Pumping energy costs are a water production operating cost for stand-alone and dual-purpose desalination plants. The result is a near waste-free deployment, since the cold thermal energy is reused (twice), costs associated with pumping energy are shared, and temperature differences between ocean depths and the sea surface can be used to both cool the servers and produce fresh water for the urban water supply.

Resource Flow Parameter	Conventional RO Plant	DOW Co-Located RO Plant
Feed Water Temperature	Ambient surface seawater	Pre-chilled via cooling loop
RO Feed Pressure Requirement	Higher (warm feed)	Lower (cold feed)
Specific Energy Consumption	Higher baseline	Reduced via pre-cooling
Membrane Fouling Frequency	Higher (surface water)	Lower (deep ocean source)
Pump Energy Attribution	100% to water production	Shared with cooling function
Water Production Cost	Higher operational cost	Lower operational cost

Table 2: Circular Resource Flow Efficiency Indicators in DOW Co-Location [4]

3. Environmental Impact and Discharge Management

In addition to technical and economic objectives, the selection of the DOW co-location model also entails environmental objectives. The extraction of cold, nutrient-rich seawater from depth and the discharge of thermal and chemical waste products associated with cooling and desalination into the coastal ocean need to be carefully managed in a systematic manner. This deals with the two main types of possible environmental harm (the physical and ecological impacts of thermal and saline plumes and the biogeochemical impacts of nutrient-laden discharges) and indicates how these can be reduced and ideally countered by good design and appropriate monitoring of these potential environmental side effects.

3.1 Thermal and Saline Plume Modeling

The overall environmental impact of a DOW-based co-location system will depend to a large extent on how the discharge stream is handled. Pumping cold nutrient-rich seawater from depth through heat exchange and desalination and then discharging it back into the coastal ocean will change the temperature and chemistry of the water column. If unabated, thermal and salinity anomalies can penetrate deep into the interior ocean where thermoclines are destabilized. Thermoclines are the boundary between habitat domains for pelagic organisms and sessile benthic communities (e.g., coral reef ecosystems), all of which are sensitive to temperature and salinity at small scales [5].

Therefore, plume modeling is not only a regulatory requirement; it is a critical engineering tool in the design of a DOW system. Computational fluid dynamics (CFD) models and 3D hydrodynamic ocean models (e.g. Princeton Ocean Model (POM); Regional Ocean Modeling System (ROMS)) can be used to simulate the plume trajectory for site-specific oceanographic conditions such as mean current velocities, tidal cycle, coastal bathymetry and seasonal stratification. The models have been used to create maps of the expected thermal and saline anomaly field as a function of discharge volume, temperature difference, outlet geometry, and ocean conditions, and enable the rapid identification of outlet geometries that provide the best plume dilution [5].

Controlled dilution is the basis for all discharge designs: the returned seawater must reach stability with the surrounding ambient ocean in terms of thermal and salinity properties before the plume boundary comes in contact with any ecologically sensitive areas. In particular, reef communities are sensitive to thermal stress and bleaching, requiring scrutiny. In estuarine and brackish water systems, where the communities are adapted to the salinity differences, the brine discharged must be diluted to the point of no measurable effect on the salinity of the near-shore water body. Hydrodynamic modeling can be used to determine appropriate depth, distance and orientation of discharge points based on oceanographic conditions, including worst-case conditions with low currents and minimum natural dilution with upwelled ocean waters [5].

3.2 Optimal Mixing Strategies

The main environmental consequence of DOW discharge is the thermal and salinity plume. Another major environmental effect of DOW is the biogeochemical signature of DOW that is discharged. DOW contains a high load of dissolved inorganic nutrients (nitrate, phosphate and silicate)discharged.e been amassed in deep water resulting from centuries of biological pumping. If these nutrients are released into a shallow oligotrophic coastal region, high concentrations of phytoplankton blooms may be formed if other physico-chemical factors (such as light intensity, water temperature and hydrodynamic residence time) are also conducive to the development of the HAB species. Some of these HABs are associated with large-scale shellfish and fish kills and potential human health threats for those consumers of marine food resources (see reference [6]).

Optimization of mixing designs is required to avoid this, but some nutrient enrichment is sometimes still permitted where demonstrated ecologically. The basic engineering design used for this purpose is called a multiport diffuser array. This consists of a discharge manifold with multiple angled, submerged ports along a header pipe. Diffuser arrays dilute nutrients, the temperature of the discharge water, and salinity differences within minutes of the discharge. This is accomplished by routing the discharge through many smaller jets, so the nutrients are mixed and entrained with the nearby seawater, thereby achieving dilution factors many times greater than a single discharge outlet pipe alone can provide [6].

In addition to diffuser arrays, staged discharge can be controlled based on real-time environmental feedback from networks of fixed and buoy-mounted sensors, measuring temperature, salinity, dissolved oxygen, chlorophyll-a fluorescence (proxy for phytoplankton biomass), and current velocity. A SCADA (Supervisory Control and Data Acquisition) system can be employed to throttle or reroute discharge in response to early warning indicators of adverse environmental impact. This adaptive, sensor-based management of discharge allows an otherwise passive engineered safety system to be used as an environmental management tool, responding dynamically to oceanographic conditions on a timescale of hours rather than days.

When properly designed and monitored, DOW diffuser discharge can provide a controlled low-level source of nutrient enrichment, which increases near-shore marine productivity without exceeding bloom-threshold concentrations. Real-time monitoring used with diffuser discharge can be aimed at optimizing zooplankton biomass, higher-trophic-level productivity, or the productivity of artisanal fisheries. Thus a discharge stream that would otherwise be an environmental liability to the ecology can be transformed into an ecologically based product, suggesting that DOW co-location infrastructure can deliver economic, social, and environmental co-benefits.

4. Socio-Technical Governance and Equitable Water Distribution

Technical and economic feasibility are necessary but not sufficient conditions for the success of a DOW co-location project over time. The long history of large-scale resource infrastructure and generation in the Global South shows yet again that technically feasible and economically viable projects threaten to fail or produce long-term social conflict when they are imposed from above rather than collaboratively produced. This section will discuss the governance architecture that may be required to transform the DOW co-location model from a commercial infrastructure asset into a true tool of water sovereignty. Institutional arrangements will be discussed to develop a social license to operate and contractual arrangements that formalize the equity of water distribution as law.

4.1 The Social License to Operate

The fact that a large-scale natural resource development project is technically and financially feasible does not necessarily guarantee its medium- and longer-term success in a developing country context. Numerous extractive and utility sector infrastructure projects in the Global South illustrate that technical feasibility does not guarantee project implementation in the absence of adequate consultation , and can take many years to arrive at first production. In many coastal developing countries where freshwater is already rare and where populations are suspicious of capital investment not coming from within the country, public tolerance for new large-scale disruptive engineering is already low. A DOW co-location system will therefore be exposed to substantial operational and reputational risk if local populations consider that it exists mainly to serve the computational needs of global technology markets at their expense [7].

Social license is not a single document signed at project approval or commencement of operations, but a continuing process of legitimacy that should be continuously earned and reconfirmed through meaningful efforts in transparent and real benefit sharing, along with meaningful community engagement and participation in project governance from the onset until project closure. For co-location projects that allow for DOWs, this means giving the communities that are affected real agency in the allocation of resources, as opposed to simply being beneficiaries of CSR programs [7].

To operationalize this commitment, local water governance committees, including local water users from underserved beneficiary households and smallholder agricultural users, women's groups, and indigenous peoples if relevant, should be empowered to make substantive decisions regarding the allocation of the desalination plant's production, monitor the data, and activate grievance mechanisms in response to non-compliance with commitments made to local communities. Project-level joint oversight boards of community representatives, national regulatory authorities and project operators should periodically review operational data, renegotiate allocations as community water demand patterns evolve, and help resolve disputes. Such tiered governance mechanisms take aspirational commitments and turn them into actionable, accountable institutional arrangements and lay the groundwork for securing a genuine social license.

In situations involving colonial and post-colonial histories of resource extraction and mistrust in institutions, social license must also be earned through substantive, early, visible and concrete benefits to the host community prior to commissioning, such as guaranteed priority hiring of local labor, investment in household water distribution infrastructure, and regular public reporting of water production and water distribution compliance metrics to show communities empirically that water sovereignty commitments are met in practice [7].

4.2 Resource Offtake Agreements

The ROA is the key binding policy mechanism through which the social license to operate is delivered. The ROA is a long-term legally binding contract and regulatory instrument between the project proponent, the host government and community governance representatives, setting out the percentage of desalinated water production that will be supplied at reduced cost to currently unserved neighboring communities. Because ROAs are legally binding obligations that run with the project through transfer of ownership, operational control and national government, they are a far more credible and effective means of ensuring long-term water security for the community than voluntary corporate social responsibility obligations that companies could renegotiate or withdraw under business pressure.

The design of the ROA must be such that the level of water allocated to the community approximately matches the projected shortfall in water supply when best estimates of real water demand per household, projected population growth, and the productive output of the desalination plant under the design contingencies for which it is designed and constructed are taken into account. The subsidized tariff must reflect levels of water affordability for poor households as defined in generally accepted definitions of water affordability, and the cost recovery mechanism must be financially sustainable for the operating concessionaire [8].

ROAs should also provide for periodic (potentially five-yearly) renegotiation of allocation volumes and tariffs to reflect factors such as the changing size of the community provided with water, local availability of alternative supply, and project financial performance. Provision for periodic renegotiation is required to ensure that ROAs do not lock parties into inequitable arrangements for the long operating life of major water infrastructure. Escalation clauses linking the percentages of community allocation to revenues from the project may address the operator's concerns by allowing the community an upside in the case of commercial success.

Independent audits of ROA compliance with unrestricted access to production records, meter data and financial records should be provided. Adequate penalties for non-compliance, including fines, community reparation payments and other regulatory penalties up to and including suspension of operating licenses, need to be specific, credible and severe enough to deter non-compliance. With codification, enforcement and audit, the ROA turns the DOW co-location project from a dual-use infrastructure asset into a legal institution of water sovereignty that gives the community a defined, proportionate share of locally produced clean water that is as secure and enforceable as any commercial contract.

5. AI-Enabled Nexus Optimization

The four subsystems of the DOW co-location model (ocean intake, data center cooling loop, desalination plant, and community water distribution system) are not simply linear subsystems. The subsystems form a socio-technical system where each subsystem depends on others and is, in turn, affected by the others depending on when decisions are made and by time-varying and stochastic demand and supply constraints. An optimization capability beyond that of any human operator or a rules-based system is thus required for nexus management. We describe here the AI-based hierarchical control architecture proposed for its optimization and design principle of hardcoded community water welfare priority so that the optimization layer becomes a mechanism to promote human development rather than commercial efficiency.

5.1 A Complex Dynamic System with Multiple Domains

All four subsystems of the DOW co-location nexus have their own independent, periodic and partially stochastic patterns of demand and supply, which do not necessarily correlate with each other. Water demand for community use has both a diurnal pattern that includes residential and agricultural use of irrigation, and is influenced by rainfall and population growth. The data center's computational load varies because of the global volume of digital information transferred, the time zone, or, temporarily, the computational load associated with training large AI models, which can lead to power draws that are multiples of the average for multiple hours. In addition, oceanographic conditions near the intake, such as water temperature, current speed, sea state, or ambient salinity, are driven by meteorological forcing, tidal cycles, and interannual climate variations such as El Niño/ENSO events [8].

All of these three independently and temporally varying signals represent a multidimensional nonlinear dynamic optimization problem. At any instant of time, the total cold seawater output of the intake pumping station must be distributed between cooling and desalination applications considering the data center thermal load, community water demand and the environmental discharge per the ecological monitoring system constraints. The complexity and high dimensionality of the problem due to continuous state variables, such as thermal loads, flow rates, membrane pressure, pump speed, levels in community reservoirs, sensor readings from the ocean in real time, etc., make it impossible for human operators and conventional rule-based control systems to optimize the operations in a real-time context. Hence, an AI optimization layer becomes structurally imperative [8].

Control Layer	Timescale	Inputs	Actions	Trained Environment
Real-Time Hardware Layer	Seconds to minutes	Live sensor data (temperature, pressure, flow)	Pump speed, valve position, feed pressure	Digital twin simulation
Short-Horizon Scheduling Layer	Hours to days	Demand forecasts, weather, workload schedules	Operational scheduling, reservoir pre-fill	Model predictive control + DRL
Long-Term Adaptive Planning Layer	Months to years	Demand trend data, capacity utilization	Capital investment recommendations	Continuous live system learning

Table 3: Hierarchical Multi-Agent Reinforcement Learning Control Architecture [9]

5.2 The Architecture of the AI Optimization Layer

The desired nexus optimization framework can be constructed using a hierarchical MARL structure. The architecture is composed of three levels, operating at different timescales: level 1 for real-time hardware control, level 2 for short-horizon scheduling, and level 3 for long-term adaptive planning. This hierarchy reflects the different timescales on which the systems variables evolve and the different levels of operational authority that can be exercised on them.

At the real-time layer, a grid of sensor and actuator control agents runs in seconds to minutes. The controllers control variable-speed intake pumps, heat exchanger bypass valves, RO feed pressure controllers, and flow controllers in the distribution network. Based on these observations, they take actions to maintain a given operating envelope for each of the subsystems, i.e., server inlet temperature, membrane differential pressure, community reservoir fill levels, ocean intake temperature, etc. The agents were trained on a high-fidelity digital twin simulation of the physical infrastructure, enabling the generation of effective control policies without requiring any closed-loop experimentation on the physical systems [9].

At a short-horizon scheduling level (on the order of hours to days), an optimizer uses community water demand forecasts (based on water consumption history, weather forecasts, calendar events, etc.), as well as data center workload and regional numerical ocean condition forecasts, to create operating schedules that pre-position the system for demand, following model predictive control and deep reinforcement learning methods. Where a peak demand for water in a community is anticipated to coincide with a time when the AI training workload is heavy, the scheduling agent might fill the community distribution reservoirs during an overnight period of low computational demand, devoting the full cold seawater throughput to the desalination process prior to the peak period [9].

5.3 Dynamic Priority Arbitration and Welfare Protection

The AI optimization layer's main ethical and technical role is to adjudicate between the demands of the data center and the local community when both cannot be met from the available DOW throughput. This should be rare in a well-dimensioned system , and only happen during maintenance for cooling equipment, on extreme weather days when the data center intake is capped, or when simultaneous demand peaks exceed the combined design basis.

First, community water demand up to the project's ROA is a hard constraint. In contrast to a soft constraint, where the system is penalized if the expected output deviates from the community water demand, the AI system is architecturally prohibited from reducing community water allocation below the project's ROA threshold regardless of its computational cost. To do this, its optimization agent must give concessions on the data center side, like batching workloads, throttling non-latency-sensitive AI inference jobs, or turning on additional refrigeration from conventional cooling, while preserving the full desalination output to the community [8]. This priority order is hard-coded into the reinforcement learning reward function and verified with formal constraint verification techniques, assuring that the optimization layer's primary goal is human welfare and its secondary goal is commercial performance.

The three layers of reinforcement learning agents re-optimize action policies based on the historical system experience amassed while the system is running. The long-term adaptive planning agent periodically monitors structural parameters (like the flow capacity of intakes, the capacity of desalination plants, and an annual forecast of operation loads) against observed demand patterns. The agent issues recommendations for capital investment when structural limits of the system are near. This enables the AI optimization layer to improve across the operational life of the system, leading to the system becoming a strong long-term infrastructure asset as both a digital and humanitarian capability.

6. Conclusion

Finally, co-locating DOW cooling with desalination technology platforms, not just as an engineering trick, but as a way to realize the most radical yet realistic descriptions of how AI infrastructure can work in the developing world, resolves the long-standing irreducible trade-off between modernization of the digital economy and water security in the developing world. DOW solves a problem we've had in the developing world for many decades: that data centers are competitors for limited local water resources. They can be a co-producer of potable water at co-located

DOW facilities. The DOW model can successfully cross-subsidize capital and operating costs to provide economically feasible water sovereignty in the developing world, where conventional desalination cannot.

Thermal and saline mixtures, flushing with multiport diffuser arrays, and practical sensor-mediated adaptive control can maintain or at least not degrade the coastal marine ecosystem and in appropriately designed schemes may provide a managed co-benefit for near-shore marine productivity. Resource Offtake Agreements and the multi-level community governance and independent auditing required to enforce them can provide hard institutional commitments that benefit all parties and ensure equitable access to water that lasts through changes of ownership, government, and commercial circumstances. This AI optimization layer's social contract prioritizes the water needs of the community but is implemented at the real-time machine-control level, with millisecond-scale decisions made by pumps, valves, and flow controllers in the continuous real-time operation of water networks in the real world.

These four dimensions can be considered as a single, reproducible template. Water-stressed coastal communities such as small island developing states in the Pacific and rapidly urbanizing seaboards in sub-Saharan Africa and South Asia need no longer choose between the intelligence revolution and their most precious resource. With proper design and governance architecture, AI infrastructure can deliver them both in equal and sovereign measure.

REFERENCES

1. David Mytton, "Data center water consumption," in NPJ, 2021. [Online]. Available: https://www.nature.com/articles/s41545-021-00101-w
2. Julian David Hunt, et al., "Deep seawater cooling and desalination: Combining seawater air conditioning and desalination," in Science Direct, Sustainable Cities and Society, Volume 74, November 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2210670721005333
3. Aminath Saadha, et al., "Techno-Economic Analysis of Combined Onshore Ocean Thermal Energy Conversion Technology and Seawater Air Conditioning in Small Island Developing States," Solar Energy, vol. 82, no. 6, pp. 520–527, 2008. [Online]. Available: https://www.mdpi.com/2071-1050/17/10/4724
4. Noreddine Ghaffour, et al., "Technical review and evaluation of the economics of water desalination: Current and future challenges for better water supply sustainability," in Science Direct, Desalination, vol. 309, pp. 197–207, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0011916412005723
5. R. A. Falconer, et al., "Hydraulic and Environmental Modelling: Coastal Waters," in Proceedings of the Second International Conference on Hydraulic and Environmental Modelling of Coastal, Estuarine and River Waters. [Online]. Available: https://api.pageplace.de/preview/DT0400.9780429826139_A38913677/preview-9780429826139_A38913677.pdf
6. Donald M. Anderson, et al., "Progress in understanding harmful algal blooms: Paradigm shifts and new technologies for research, monitoring, and management," Annual Review of Marine Science, vol. 4, pp. 143–176, 2012. [Online]. Available: https://www.annualreviews.org/content/journals/10.1146/annurev-marine-120308-081121
7. Johannes Glückler and Denise Gutiérrez, "Social license to operate: an institutional critique and research framework," Resources Policy, Volume 107, August 2025, 105657. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0301420725001990#:~:text=Highlights,understand%20place%2Dspecific%20SLO%20dynamics.
8. WHO and UNICEF Joint Monitoring Programme, "Progress on household drinking water, sanitation and hygiene 2000–2020: Five years into the SDGs," World Health Organization, 2021. [Online]. Available: https://www.who.int/publications/i/item/9789240030848
9. Volodymyr Mnih, et al., "Human-level control through deep reinforcement learning," 2015. [Online]. Available: https://user.ceng.metu.edu.tr/~emre/resources/courses/AdvancedDL_Spring2017/DQN_Muhammed.pdf
10. Albin Heimerson, et al., "Adaptive Control of Data Center Cooling using Deep Reinforcement Learning," in Research Gate, 2022 IEEE International Conference on Autonomic Computing and Self-Organizing Systems Companion (ACSOS-C) 2022. [Online]. Available: https://www.researchgate.net/publication/365137031