

Breeze First, Seawater Second: Hybrid Cooling for an Above-Water Offshore Data Center on San Francisco Bay Within California's Discharge Limit

Vishnu Sai Sharan Ankathi · Independent research · ORCID 0009-0003-9824-3632 ·
vishnusaisharan.a@gmail.com · Preprint v1, October 2026

Abstract

Coastal sites offer data centers two free heat sinks: cool marine air and seawater. Existing Bay-adjacent concepts pump water continuously, which has already drawn regulatory opposition over thermal discharge and marine-life entrainment. This screening-level feasibility study evaluates a hybrid, above-water design for a 270 kW server room on a pile-supported San Francisco Bay pier. An indirect air-to-air heat exchanger uses the marine-layer breeze as the primary heat sink, an air-to-seawater coil is pumped only when the breeze cannot hold the 21 °C supply-air setpoint, and a small chiller trims the remainder. A room-level CFD model (Ansys Icepak) sets the airflow and hot-aisle return temperature (33.25 °C); an hourly ϵ -NTU model of the two-stage heat-exchanger chain is then driven by three years of hourly Bay air data (2022–2024) and measured Bay water temperature; fan, pump and chiller power are then added to estimate annual energy and mechanical PUE against two baselines; finally, the air-to-air exchanger is sized from first principles to replace the assumed pressure drop.

The breeze alone holds the setpoint in 84.8% of hours, seawater is needed in 15.1%, and the chiller in 0.14%. With an assumed compact exchanger (200 Pa per side) the hybrid uses 245–275 MWh of cooling electricity per year (mechanical PUE 1.10–1.12), about 63–67% less than an air-cooled chiller plant. California's Thermal Plan prohibits new cooling-water discharges into San Francisco Bay more than 2.2 K (4 °F) above the natural water temperature. A fixed-flow seawater loop breaks this limit, so the hybrid and a seawater-first baseline are both re-modelled with a variable-speed pump that holds the rise at 2.0 K. Under that rule, seawater first pumps about 1.13 million m³ of Bay water a year and uses 182–219 MWh (PUE about 1.08), depending on pipe friction. With the air-to-air exchanger sized as a plain aluminium plate stack (3 mm gap, 1.2 m plates, 67 Pa per side), the hybrid uses 187–217 MWh (PUE 1.08–1.09): about the same energy, while pumping about 27,000 m³, about 41 times less Bay water and heat. The stacks are about eight times larger than first drawn, and the result depends on an exchanger effectiveness near 0.75; at 0.60, typical of some commercial cores, pump runtime more than triples. A commercial counterflow core checked in its maker's selection tool reaches 81–85%, so 0.75 is a conservative assumption. Raising the supply setpoint to 24 °C, within ASHRAE's recommended range, cuts pump runtime to about 4.4% of hours; room CFD confirms that, with full cold-aisle containment, rack inlets stay at or below 25.5 °C at that setpoint. Treating the room as a repeatable module, the percentages carry over to 1–10 MW facilities, where exchanger space, exhaust recirculation (each 1 K of intake warming adds 7–11 points of pump runtime) and total Bay withdrawals become the binding constraints.

1. Introduction

Cooling accounts for a large share of data center energy use, and the burden is growing as AI and cloud workloads raise power density. At the same time, land near major compute hubs is scarce and freshwater cooling is under scrutiny as data centers compete with communities for water during droughts.

The San Francisco Bay Area coastline offers a climate resource that existing water-cooled concepts in the region leave unused: a persistent marine-layer breeze that keeps coastal air cool through much of the year. Nautilus Data Technologies has shown that water-adjacent siting is commercially viable in the region, with a barge-mounted facility in Stockton operating since 2021, but its design and its earlier, rejected Alameda pier proposal rely on continuous pumped-water cooling. That places the full thermal-discharge and intake burden on the Bay.

This work asks how much of that pumping can be displaced by the air already moving past the site, and at what cost. We evaluate a hybrid design in which marine air, drawn through an indirect heat exchanger, handles the baseline load and pumped seawater covers only the residual. Rather than relying on benchmark PUE figures, the analysis is built bottom-up: a room CFD model fixes airflow and temperatures, an hourly heat-exchanger model follows the cooling chain from outside air to the racks, and a parasitic-power model accounts for every fan and pump. The main contribution is an hourly, discharge-constrained comparison of an air-first hybrid against seawater-first cooling for a Bay-adjacent data-center module, including fan, pump and chiller power and a room-level airflow basis. Without the discharge limit, seawater first uses the least energy; with it, the ranking changes.

2. Background and Positioning

2.1 Free-air and economizer cooling

Free-air and economizer cooling are mature strategies. Meta's Prineville, Oregon facility reports a PUE of about 1.07 [1], against an industry average nearer 1.55 [2]. Indirect air-side economizers, in which outside air cools the data-hall air through a heat exchanger without entering the IT space, avoid the humidity and particulate risks of direct free-air cooling and are the architecture adopted here.

2.2 Water-cooled and marine-adjacent data centers

Nautilus Data Technologies operates a barge-mounted, water-cooled facility at the Port of Stockton (about 6.5 MW critical IT load, company-reported PUE 1.15, roughly 4,500 GPM drawn from the San Joaquin River) [3], [4]. Nautilus separately proposed a pier-mounted facility at Alameda Point that would have drawn and discharged about 10,000 GPM continuously; the Alameda City Council rejected it unanimously in 2019 after environmental objections centered on thermal discharge and marine-life entrainment [5], [6]. Water-cooled Bay-adjacent data centers are therefore technically viable, but continuous once-through seawater use faces a real regulatory ceiling in this region.

Seawater cooling is also in commercial use elsewhere. Start Campus's SINES campus in Portugal uses Nautilus seawater cooling [7], and a 2024 study evaluated air-only and hybrid air/liquid server cooling with seawater as the heat sink [8]; there, "hybrid" refers to the mix of air- and liquid-cooled IT equipment, not to an outdoor-air stage. Google's Hamina facility in Finland mixes its warm seawater outlet with additional

cold seawater so that it returns close to sea temperature [9]. Deep-water studies report large savings from seawater-only cooling, for example up to 78% for a 700 m intake in the Caspian Sea [10], and hybrid schemes that switch between free water cooling and a chiller, such as deep dam water [11], still have no outdoor-air stage. Reviews of offshore and floating data centers report seawater-cooled PUEs of 1.07–1.15 [12], and subsea modules such as HiCloud’s wind-powered installation off Shanghai reject heat directly to the sea [13]. At grid scale, moving from water-cooled to air-cooled systems trades direct water use for electricity and indirect water use [14]. None of these studies uses outdoor air as the primary heat sink or applies a discharge-temperature limit to the seawater loop.

2.3 Offshore corrosion engineering

The offshore oil-and-gas industry provides a quantitative literature on marine corrosion that has not been applied to above-water data centers. Corrosion protection has been estimated at up to 10% of platform capital cost [15] and corrosion-related maintenance at 25–40% of operating cost (industry estimates). Splash-zone protection is the most expensive per unit area: copper-nickel sheathing on Shell’s Eureka platform was reported at about \$340/m² installed [16]. Cathodic protection is designed to the full structural life (20–40 years, DNV-RP-B401 [17]), and glass-fiber-reinforced polymer (GFRP) has been reported to cut offshore topside life-cycle cost by up to 26% over 50 years [18].

2.4 Gap addressed

No identified study models a hybrid marine-air and seawater architecture in which air carries the baseline load and water covers only the residual, and none quantifies the resulting trade-off between cooling energy and seawater use. The closest work [8] pairs seawater with air- and liquid-cooled servers but has no outdoor-air stage. Google’s patented approaches (e.g., US 8,853,872 [19]) and pier-mounted heat-exchanger concepts (US 6,736,191 [20]) cover related elements, so the contribution claimed here is narrower: an hourly, discharge-constrained comparison of an indirect air-first hybrid against seawater-first cooling for an above-water Bay data-center module, using measured Bay air and water data, fan and pump power, and a room-level airflow basis, plus a first-pass duty-cycle view of seawater-loop maintenance.

3. Reference System and Methods

3.1 Site data

Hourly air temperature, relative humidity and dew point for a Bay coastal point (37.79 °N, 122.35 °W) were taken from the Open-Meteo historical archive [21], [22] for 2022–2024 (26,304 hours). Hourly Bay water temperature came from NOAA CO-OPS station 9414290 (San Francisco) [23]. Water records end in February 2024; the missing hours are filled with month-by-hour averages of the measured data and flagged, and worst-case days are selected from measured hours only. Mean air temperature over the period was 13.5 °C and mean relative humidity 78.7%.

3.2 Reference room and cooling architecture

The reference room measures 5.79 × 4.965 × 2.5 m and holds 18 standard racks (600 × 1000 mm) in two rows of nine, with a contained 1.2 m cold aisle and 0.88 m hot aisles. This is the largest count that fits with

realistic service clearances. At 15 kW per rack, a typical density for air-cooled enterprise and cloud halls, the IT load is 270 kW. The room is treated as a repeatable module: a larger facility is built from N identical rooms, each with its own cooling paths (Section 6.3). Cooling is split into two independent paths of 135 kW, one per row, each carrying $10.06 \text{ m}^3/\text{s}$ of room air (42,650 CFM total at an 11–12 K air temperature rise).

Each path runs outside air and sealed room air through separate passages of the same exchanger (Figure 1). Outside air enters through three filter banks and an intake plenum, passes through the Stage 1 air-to-air exchanger, and is exhausted; it never reaches the IT space. Room air leaves the hot aisle, crosses Stage 1, then passes an air-to-seawater coil (Stage 2) and returns to the cold aisle. Sizing a single exchanger for the full load produced impractical faces (about 2.8 m square), so each path is split into three parallel sub-units with $1.15 \times 1.15 \text{ m}$ faces, which fit the plenum and add redundancy. Section 4.6 shows these panels must grow into taller stacks to reach the design effectiveness.




Figure 1. Concept CAD of the pier-mounted facility (SolidWorks): intake filter banks and mixing plenum, two Stage 1 rows of three $1.2 \times 1.2 \times 2.7 \text{ m}$ plate-exchanger cores with their plenums, two air-to-seawater coils with Bay-water pipes through the deck, and the sealed server room with its return and supply ducting. One cooling path per row.

3.3 Thermal criteria

Supply air is controlled to 21°C , inside the ASHRAE 2021 recommended range of $18\text{--}27^\circ\text{C}$ for Class A2 equipment [24]; 24°C is evaluated as a sensitivity case. Direct free-air cooling was screened against both the recommended envelope ($18\text{--}27^\circ\text{C}$, $5.5\text{--}15^\circ\text{C}$ dew point, $\leq 60\%$ RH) and the wider A2 allowable envelope. The indirect architecture is assessed on dry-bulb temperature only, because humidity does not cross the exchanger.

3.4 Room CFD model

A steady-state Ansys Icepak (AEDT 2026 R1) model of the room represents each rack as a recirculating object that adds 15 kW and moves its share of the airflow, with 21 °C supply at 10.06 m³/s per path, turbulent flow and gravity, on a stair-stepped mesh of roughly 96,000 cells (100 × 100 × 80 mm maximum element). The run converged in about 145 iterations (continuity residual ~10⁻³, energy ~10⁻⁷), and the airflow through the room was 42,564 CFM against the 42,650 CFM design (0.2% low). The cold aisle held 21 °C and hot-aisle air peaked at 33.25 °C, against a hand-calculated 32.1 °C (Figure 2). Some hot air leaks around the row ends; Section 4.8 tests containment against this. The 33.25 °C return is used as the inlet condition for the heat-exchanger model, which makes that model slightly conservative (the implied heat removal is about 292 kW against the 270 kW IT load). With full containment (Section 4.8) the mean return is 32.45 °C; adding the 0.8–0.95 K of room-fan heat that the CFD omits gives about 33.3 °C, consistent with the 33.25 °C used.

CFD check: 21 °C supply keeps 270 kW of racks in spec

Ansys Icepak · 18 racks × 15 kW · hot and cold aisles stay separated




Figure 2. CFD Test 1: air temperature on a horizontal slice through the room. The 21 °C cold aisle and ~33 °C hot aisles stay separated.

3.5 Hourly heat-exchanger chain

For each hour and each path, the model follows the room air through both stages. Stage 1 is a cross-flow exchanger with both fluids unmixed. Its UA is fixed by a design effectiveness of 0.75 at equal outside and room airflow, and the effectiveness then follows the NTU relation as the outside-air flow changes [25]:

$$\varepsilon = 1 - \exp\left\{\left(1/C_r\right) \cdot NTU^{0.22} \cdot \left[\exp(-C_r \cdot NTU^{0.78}) - 1\right]\right\}$$

Stage 2 is a counterflow air-to-seawater coil with design effectiveness 0.80 at 3.25 kg/s of seawater per path. Control is sequential: Stage 1 runs first and is modulated so it never overcools; if the air leaving Stage 1 is still above setpoint the seawater pump starts; any remaining heat goes to a chiller (COP 5). Stage 1 is

bypassed in the rare hours when outside air is warmer than the return air. With these values, the breeze alone meets 21 °C whenever outside air is at or below 16.9 °C.

3.6 Parasitic power, baselines and mechanical PUE

Fans and pumps are sized from pressure drop and flow, $P = \dot{V} \cdot \Delta p / \eta$. Room-air fans run continuously; outdoor-air fans are variable-speed and run only as fast as needed to reach setpoint (power scales with speed cubed, minimum 20% speed); the seawater pump runs only in Stage 2 hours. Table 1 lists the design values. They are typical engineering values rather than vendor data and are varied in Section 4.6. Mechanical PUE is defined as (IT load + cooling electricity) / IT load; UPS and lighting losses are excluded because they are the same for every option.

Table 1. Design assumptions for the parasitic-power model (per cooling path).

Item	Value	Basis
Room-air loop pressure drop	680 Pa	Ducts 250, Stage 1 200, Stage 2 coil 120, containment/filter 110
Outdoor-air loop pressure drop (design flow)	550 Pa	MERV 13 filters 150, intake 100, Stage 1 200, exhaust 100
Fan + motor + VFD efficiency	0.60	Typical packaged air handler
Seawater pump head / efficiency	160 kPa / 0.60	~5 m lift to deck, coil, strainer, piping
Resulting design power	11.4 / 9.2 / 0.85 kW	Room fans / outdoor fans / pump
Trim chiller COP (hybrid, seawater-first)	5.0	Water-cooled trim unit
Air-cooled chiller COP (baseline)	4.0 (3–6 tested)	Including condenser fans

Two baselines use the same room, the same room airflow and the same data. Chiller only replaces both exchangers with a chilled-water coil (150 Pa) served by an air-cooled chiller every hour. Seawater first removes the air stage and pumps Bay water through the Stage 2 coil every hour, with chiller trim; it represents the conventional way of using the Bay. A fourth case, hybrid with coil bypass, lets room air skip the dry seawater coil when the pump is off.

3.7 Sizing the Stage 1 exchanger

The 200 Pa assumed above is typical of compact commercial units. To test it, Stage 1 is sized as a cross-flow stack of flat 0.15 mm aluminium plates with alternating room-air and outdoor-air channels, three cores per path, each carrying 3.35 m³/s on both sides. Reaching $\epsilon = 0.75$ at equal flows requires $NTU \approx 5.05$, or $UA \approx 20.1$ kW/K per core. Channel Reynolds numbers are 150–1,050, so the flow is laminar: heat transfer uses the fully developed parallel-plate value $Nu = 7.54$ ($h = Nu \cdot k / Dh$, $Dh = 2 \times \text{gap}$), independent of velocity, and pressure drop uses $f \cdot Re = 96$ [26] plus one velocity head for entry and exit. With equal h on both sides the required plate area is $A = 2 \cdot UA / h$. Gaps of 3–8 mm and plate sizes of 0.6–1.5 m were swept, and the resulting pressure drops were fed back into the energy model.

3.8 Key assumptions and applicability

Table 2 collects the assumptions behind the headline results and what the study does not cover. The results apply to a single 270 kW module at the stated conditions; facility-scale behaviour needs the separate analyses listed in the last row.

Table 2. Key assumptions and scope.

Item	Value used	Basis
IT load	270 kW per module: 18 racks $\times$ 15 kW, two 135 kW paths	Design basis
Room airflow and return	10.06 m ³ /s per path; return 33.25 °C	CFD Test 1 (Section 3.4), checked by Test 2B
Supply setpoint	21 °C base; 24 °C case	ASHRAE recommended 18–27 °C [24]
Stage 1 exchanger	Crossflow, $\epsilon = 0.75$ at equal flows; 67 Pa (sized) or 200 Pa (assumed)	Sections 3.7, 4.6; commercial core 81–85%
Stage 2 coil	Counterflow, $\epsilon = 0.80$ at 3.25 kg/s per path; fixed UA	Design assumption, not vendor data
Fans, pump, chillers	Fan η 0.60; pump 160 kPa, η 0.60; trim chiller COP 5; air-cooled baseline COP 4	Table 1
Seawater rise	≤ 2.0 K modelled (rule 2.22 K)	Thermal Plan [27]; screening constraint only
Weather and water data	Hourly air 2022–2024; NOAA water to Feb 2024, gaps filled by climatology	Open-Meteo/ERA5 [21], [22]; NOAA [23]
Not included	Capital cost, site-scale CFD and exhaust recirculation, plume and mixing-zone modelling, entrainment, permit determination, rack internals	Sections 6.3–6.5

4. Results

4.1 Direct free-air cooling is ruled out by humidity

Only 3.9% of hours fall inside ASHRAE's recommended envelope for direct free-air cooling, and 40.0% inside the wider A2 allowable envelope. Temperature is rarely the problem at this site; humidity is. This confirms the choice of an indirect architecture, which also keeps salt-laden air away from the IT equipment.

4.2 Hourly duty cycle

With the two-stage chain at 21 °C supply, the breeze alone holds the setpoint in 84.8% of hours, Bay water is added in 15.1%, and the chiller is needed in 0.14% of hours, with a peak of about 17 kW thermal. Summer and early autumn, not winter, are the difficult season (Figure 3). In September the breeze alone covers 60% of hours, because coastal air temperature does not fall in summer even when the sea breeze strengthens, while Bay water stays cool (Figure 4).




Figure 3. Share of hours by cooling stage, month by month, 2022-2024.




Figure 4. Hottest day with measured Bay water (6 September 2022): the breeze pre-cools the room air and the seawater coil brings it to 21 °C.

4.3 Annual energy and mechanical PUE

Table 3 and Figure 5 give annual cooling electricity with the assumed 200 Pa Stage 1 exchanger (Section 4.6 replaces this with a sized design). The hybrid needs 275 MWh per year (mechanical PUE 1.12), or 245 MWh (PUE 1.10) with the coil bypass, 63–67% less than the chiller-only plant at 741 MWh (PUE 1.31). Against a more efficient chiller baseline (COP 6, 544 MWh) the saving is still about 50–55%. Once the chiller is removed, fans dominate: room-air fans alone account for about 70% of the hybrid’s cooling energy.

The seawater-first design uses the least energy, 157 MWh per year (PUE 1.07). The seawater pump draws only 0.85 kW per path, while pulling outdoor air through filters and an air-to-air exchanger takes about 9 kW per path at design flow and adds 200 Pa to the room-air loop. On energy alone, the breeze does not beat the Bay. This comparison, and Sections 4.4–4.6, use a fixed seawater flow of 3.25 kg/s per path, which returns water up to 11 K warmer than the Bay and would not be permitted. Section 4.7 repeats the comparison with a compliant loop, and the energy ranking changes.

Table 3. Annual cooling electricity and Bay water use, 270 kW IT load, 21 °C supply.

Option	Cooling electricity (MWh/yr)	Mech. PUE	Pump runtime	Seawater pumped (m ³ /yr)
Chiller only	741	1.31	0%	0
Seawater first	157	1.07	100%	200,000
Hybrid (breeze first)	275	1.12	15.2%	30,500
Hybrid + coil bypass	245	1.10	15.2%	30,500




Figure 5. Annual cooling electricity by component. The chiller dominates the baseline; fans dominate every free-cooling option.

4.4 Bay water use and heat discharge

The picture reverses when Bay impact is considered (Figure 6). The hybrid pumps about 30,000 m³ of seawater per year against 200,000 m³ for seawater first, a reduction of about 85%, and the pump runs in 15% of hours instead of all of them. Heat returned to the Bay falls from essentially the full IT load (about 2.4 GWh thermal per year) to roughly 0.06 GWh, about forty times less. Priced at \$0.125/kWh, the hybrid's extra 118 MWh relative to seawater first costs about \$14,700 per year, or roughly \$0.09 per cubic metre of seawater not pumped. With a discharge-compliant loop the difference grows to about 41 times (Section 4.7).




Figure 6. Bay water pumped through the coils per year.

4.5 Sensitivity

Table 4 varies the design choices and the uncertain assumptions one at a time. Raising the supply setpoint to 24 °C is the most effective lever: pump runtime drops to 4.4% of hours, seawater use to about 8,700 m³ per year and PUE to 1.10, with no chiller hours. This case assumes the 12.25 K rack temperature rise from CFD Test 1 is unchanged, so return air rises to 36.25 °C; CFD Test 3 (Section 4.8) confirms both: rack inlets stay at or below 25.5 °C, and the return, once fan heat is added, matches the assumed 36.25 °C. A higher outside-air flow ratio (1.5) or a more effective Stage 1 exchanger (0.85) both cut pump runtime to about 7%. Fan pressure drop and efficiency are the largest energy uncertainties: ±30% on all pressure drops moves PUE between 1.08 and 1.15, but the hybrid remains well below the chiller-only plant in every case.

Table 4. One-at-a-time sensitivity of the hybrid design (baseline: Stage 1 ϵ 0.75, outside/room airflow 1.0, 21 °C supply).

Case	Cooling electricity (MWh/yr)	Mech. PUE	Pump runtime
Baseline	275	1.12	15.2%
Supply 24 °C	243	1.10	4.4%
Outside-air ratio 1.5	249	1.11	6.7%
Stage 1 ϵ 0.65	309	1.13	36.9%
Stage 1 ϵ 0.85	257	1.11	7.0%
All pressure drops +30%	356	1.15	15.2%
All pressure drops -30%	193	1.08	15.2%
Fan efficiency 0.50	329	1.14	15.2%
Fan efficiency 0.70	236	1.10	15.2%
Pump head 250 kPa	276	1.12	15.2%

4.6 Stage 1 exchanger size and real pressure drop

Because laminar heat transfer depends only on the gap, the required plate area, and hence the exchanger volume, is set by the gap alone: about 12 m³ per path at 3 mm, 20 m³ at 4 mm and 45 m³ at 6 mm. Plate size only changes the shape, trading a short, wide stack with more pressure drop for a tall, narrow one with less (Figure 7). Every plain-plate design stayed under about 105 Pa per side, well below the 200 Pa assumed.




Figure 7. Stage 1 pressure drop per side versus plate size for each gap, at $\epsilon = 0.75$. The dashed line is the 200 Pa assumed in Section 4.3.

The chosen design uses a 3 mm gap and 1.2 m plates: 855 plates per core ($1,230 \text{ m}^2$), a $1.2 \times 1.2 \text{ m}$ stack 2.7 m tall, and 67 Pa per side. Six cores occupy about 8.6 m^2 of deck and 23 m^3 in total, roughly eight times the volume drawn in the original CAD, which showed $1.15 \times 1.15 \times 0.35 \text{ m}$ panels. The stacks fit on the pier but the layout must be redrawn around them. Table 5 shows the energy effect.

Table 5. Hybrid cooling electricity with the sized Stage 1 exchanger.

Stage 1 exchanger	Pressure drop	Cooling electricity (MWh/yr)	Mech. PUE	With coil bypass
Assumed compact unit	200 Pa	275	1.12	245 (1.10)
Plain plates, 3 mm, 1.5 m	105 Pa	234	1.10	204 (1.09)
Plain plates, 3 mm, 1.2 m (chosen)	67 Pa	218	1.09	188 (1.08)
Plain plates, 4 mm, 1.5 m	33 Pa	204	1.09	174 (1.07)

With the sized exchanger, the hybrid's energy premium over seawater first falls from 118 to 61 MWh per year, or from 88 to 31 MWh with the coil bypass: about \$7,600 or \$3,900 a year at \$0.125/kWh, or \$0.02–0.05 per cubic metre of seawater not pumped. Section 4.7 shows that this premium disappears once both designs meet the Bay discharge limit.

As a check on the 0.75 assumption, a commercial counterflow core (Recutech REC+120, $1425 \times 1200 \times 600 \text{ mm}$ aluminium, 3.3 mm plate spacing) was run in the manufacturer's online selection tool [28] at the design conditions, with equal mass flows of outdoor air at 13°C and room air at 33.25°C . Dry effectiveness was 84.7% at $1,500 \text{ m}^3/\text{h}$ per unit (about 23 units per path, 71–78 Pa), 81.9% at $2,500 \text{ m}^3/\text{h}$ (14 units, 148–164 Pa) and 81.0% at $3,000 \text{ m}^3/\text{h}$ (12 units, 196–218 Pa). Run through the discharge-compliant model with coil bypass, these give 174, 210 and 232 MWh per year, pump runtimes of 7.3%, 9.1% and 9.4%, and $15,000$ – $18,400 \text{ m}^3$ of Bay water, against 187 MWh, 15.2% and $27,400 \text{ m}^3$ for the sized crossflow stack.

Counterflow needs an NTU of about 3.0 for 0.75 effectiveness against about 5.0 in crossflow, so the crossflow stack assumed here is the harder case. The low-pressure option needs about twice the core volume (24 m³ per path); manifold losses and price are not included and need a vendor selection to confirm.

4.7 Bay discharge limit and robustness checks

California's Thermal Plan names San Francisco Bay as an enclosed bay and, for new discharges, prohibits thermal waste more than 4 °F (2.22 K) above the natural temperature of the receiving water [27]. With the fixed 3.25 kg/s per path used above, the seawater returns up to 11.3 K warmer than the Bay. The hybrid exceeds the limit in 26% of its pumping hours and seawater first in every hour. The Stage 2 coil was therefore re-modelled with a variable-speed pump: each hour the pump runs at the lowest flow that both reaches the air setpoint and keeps the water rise at or below 2.0 K, with the coil's conductance fixed at its original design value. Pipes are sized for the peak flow (50 kPa static lift plus friction rising with flow squared to 160 kPa at peak). In this paper, "discharge-compliant" means only that the modelled seawater temperature rise stays within 2.0 K; it is a screening constraint, not a permit determination.

The hybrid's yearly seawater use barely changes (27,400 m³ instead of 30,500 m³), because volume follows heat moved divided by the temperature rise. The hardware does change: the intake, pump and pipes must carry about 18 kg/s per path (about 280 GPM), 5.6 times the earlier design flow. Seawater first is hit much harder. It must pump about 1.13 million m³ a year, and its pump energy rises from about 15 to 60 MWh. Table 6 and Figure 8 show the result: with both designs compliant, the hybrid with the sized exchanger and coil bypass uses 187 MWh a year against 201 MWh for seawater first, and 217 MWh without the bypass. Seawater first is sensitive to the assumed pipe friction (182 MWh at a 100 kPa design head, 219 MWh at 220 kPa); the hybrid is not, because its pump runs 15% of the time. The fair reading is energy parity within about ±10%, with the hybrid withdrawing and heating about 41 times less Bay water.

Table 6. Annual cooling electricity and Bay water with both designs holding the discharge within 2 K of Bay temperature (21 °C supply).

Design (discharge-compliant)	Cooling electricity	Mech. PUE	Bay water / yr
Seawater first	201 MWh (182–219)	1.08	1,128,000 m ³
Hybrid, assumed 200 Pa exchanger	273 MWh	1.12	27,400 m ³
Hybrid, 200 Pa + coil bypass	243 MWh	1.10	27,400 m ³
Hybrid, sized 67 Pa exchanger	217 MWh	1.09	27,400 m ³
Hybrid, sized 67 Pa + coil bypass	187 MWh	1.08	27,400 m ³




Figure 8. Cooling electricity and Bay water pumped when both designs meet the 2.2 K discharge limit.

Federal Clean Water Act 316(b) rules for new facilities apply above a design intake flow of 2 million gallons per day [29]. One 270 kW module needs 0.82 MGD at peak, below the threshold; at 1 MW the hybrid needs about 3 MGD, as does seawater first, because both designs need a similar peak flow. The hybrid's regulatory advantage is therefore the volume actually withdrawn and the heat returned, not the design intake flow. A 0.5 ft/s through-screen velocity needs only about 0.24 m² of open screen area per module.

Five further checks were run. Fan heat: the room fans add 18–23 kW, so the loop must remove 288–293 kW, which the CFD-based return temperature (292.6 kW) already covers. Fouling: a 0.5 mm deposit on each outdoor-side wall raises a 3 mm gap's pressure drop from 67 to 224 Pa but adds only about 16 MWh a year; a 4 mm gap rises from 33 to 78 Pa and is the more tolerant choice for salt air. Condensation: with the room dew point at or below 12 °C (about 56% relative humidity at 21 °C supply) neither the Stage 1 plates nor the seawater coil condense; at the 15 °C upper recommended dew point the plates would condense in about 900 hours a year. Exhaust recirculation: a uniform intake rise of 1, 2 and 3 K raises pump runtime from 15% to 22%, 31% and 43% and energy from 187 to 200, 214 and 229 MWh. Commercial cores: published dry efficiencies range from about 60% [30] to about 80%; at $\epsilon = 0.60$ the breeze covers only 45% of hours and pumping rises to 97,000 m³ a year, so a core rated at 0.75 or better is a firm design requirement.

4.8 Room containment and 24 °C supply (CFD Tests 2 and 3)

Test 1 showed warm air leaking around the 195 mm row-end gaps. Test 2 first added thin adiabatic panels in those gaps (2A), then full cold-aisle containment with panels from the rack tops to the ceiling along both rows (2B). Rack-inlet temperature was sampled along each row 50 mm into the cold aisle at 1.4 m height,

and the mean return temperature was taken on both hot-aisle ceiling openings. End panels alone sealed the row ends at mid-height but let hot air spill over the racks, so inlets near the rack tops reached 22–24 °C along most of the row (Figure 10). Full containment held inlets at 21.0 °C along both rows, rising to 22.5 °C only within 0.5 m of the row ends (Table 7).

Test 3 repeated 2B with 24 °C supply. Inlets stayed at 24.0–25.5 °C, inside the 27 °C ASHRAE recommended limit (Figure 9), and the return mean was 35.45 °C in both aisles: an 11.45 K rise, the same as at 21 °C. The energy balance closes within about 1.5% ($10.06 \text{ m}^3/\text{s} \times 1.18 \text{ kg}/\text{m}^3 \times 1006 \text{ J}/\text{kg}\cdot\text{K} \times 11.45 \text{ K} \approx 137 \text{ kW}$ per path against 135 kW). Adding the 18–23 kW of room-fan heat, which the CFD does not include, gives about 12.3 K, matching the 12.25 K used in the energy model, so the 24 °C case in Section 4.5 stands. Refining the mesh from about 96,000 to 566,640 cells changed the mean return temperature by less than 0.02 K; all runs converged in about 130–150 iterations. A one-cell spike at the junction of the zero-thickness end panel and the rack face is an interpolation artifact: it disappears on the refined mesh and is excluded (Figure 10).

Table 7. Room CFD results. Rack inlets sampled 50 mm into the cold aisle at 1.4 m height along both rows; return is the mean over each hot-aisle ceiling opening. -: not measured; Test 1 and 2A were assessed on rack inlets only.

Run	Supply	Rack inlets	Return (mean)	Rise	Cells
Test 1, open aisles	21 °C	21 °C, leakage at row ends	33.25 °C (peak)	–	~96k
2A, end panels only	21 °C	22.0–24.0 °C near rack tops	–	–	~96k
2B, full containment	21 °C	21.0–22.5 °C	32.45 °C	11.45 K	~96k
Test 3, full containment	24 °C	24.0–25.5 °C	35.45 °C	11.45 K	~96k
Mesh check (Test 3)	24 °C	same pattern	35.44 °C	11.44 K	566,640




Figure 9. Air temperature on a horizontal plane at rack height: (a) end panels only, (b) full cold-aisle containment, (c) full containment at 24 °C supply, (d) the same case on a 566,640-cell mesh. All four planes are at 1.4 m, just below the 1.5 m rack tops; each panel has its own colour scale. With end panels only, warm air spilling over the racks shows as a lighter band along the cold-aisle edges (a); full containment removes it (b–d). Figure 10 gives the inlet temperatures.




Figure 10. Rack inlet temperature along row A, sampled 50 mm into the cold aisle at 1.4 m height. End panels alone leave most of the row 1–3 K above supply. Full containment holds the supply temperature except within 0.5 m of the row ends, and the 24 °C case is the 21 °C case shifted by 3 K. The refined mesh reproduces the profile. Curves are digitised from the Icepak line plots; the first 0.12 m of the ~96k-cell runs is omitted (plate-edge artifact, absent on the refined mesh).

5. Seawater-Loop Maintenance and Structural Corrosion

The platform’s legs are wetted by tides and waves whether or not the pump runs, so splash-zone protection is a fixed, site-driven cost that the cooling design does not change. For an illustrative four-legged substructure (0.5 m legs, 2 m splash zone) the exposed area is about 12.6 m², costing roughly \$13,000 for copper-nickel sheathing or about \$1,200 for a coating system before corrosion allowance. GFRP legs are a separately justified option.

The seawater intake, pump, piping and coil, by contrast, are wetted in proportion to runtime. For an illustrative \$100,000 loop package with continuous-duty maintenance at 10% of capital per year, a linear duty-cycle model gives an 85% reduction in loop maintenance at the hybrid’s 15.2% runtime. A more conservative “floor” model, cost = baseline × (0.30 + 0.70 × duty), which reflects evidence that intermittent wetting drives corrosion out of proportion to exposure time, gives about \$4,100 per year against \$10,000, a 59% reduction. At a 24 °C setpoint the floor model gives about \$3,300 (67% reduction). The equipment cost and maintenance rate are placeholders; the distinction between fixed structural and runtime-dependent loop cost is the more robust result.

6. Discussion

6.1 Once the discharge rule applies, the Bay is not the cheaper heat sink

With a fixed seawater flow, pumping Bay water every hour used the least energy, because a pump moves far less mass than a fan for the same heat. That comparison ignored the Bay discharge limit. A loop that may warm the water by only 2.2 K must move about five times more water at peak, and a design that pumps every hour then pays for it in every hour. Applying the same rule to both designs, the hybrid with a sized exchanger and coil bypass matches seawater first on energy (mechanical PUE 1.08 for both) while withdrawing and heating about 41 times less Bay water. The air stage therefore buys a large reduction in

intake volume, entrainment exposure and thermal discharge, the issues that stopped the Alameda proposal, at little or no energy cost. That holds only if the air-to-air exchanger is large, low in pressure drop and close to 0.75 effective; with a compact 200 Pa unit or a 0.60-effective core, the hybrid uses 20–55% more energy than seawater first.

6.2 Design levers

- Supply setpoint: 24 °C cuts seawater use by about 70% and lowers energy at the same time, because outdoor-air fans can run slower. It should be the default if the IT equipment allows it; CFD Test 3 shows rack inlets stay within ASHRAE limits with full containment.
- Pressure drop: room-air fans run 24/7 and dominate energy. A coil bypass when the pump is off saves about 30 MWh per year, and low-pressure-drop exchangers and short duct runs are worth more than any control refinement.
- Exchanger type: large plain-plate stacks give low fan power; compact corrugated units, typical of commercial products, fit easily but cost roughly 150–250 Pa. The choice is space against fan energy.
- Outside-air flow ratio: a larger outdoor-air flow cuts pump hours but costs fan energy in summer; the right ratio depends on how strongly the operator weights seawater use.

6.3 Scaling to 1 MW and 10 MW

Every quantity in the model is proportional to IT load: airflow, exchanger area, fan and pump power, and seawater flow all scale per kilowatt. If a facility is built from repeated 270 kW modules, the percentages found here carry over unchanged: breeze-only hours, mechanical PUE, the saving relative to a chiller plant and the reduction in seawater use. Absolute totals scale linearly, as Table 8 shows.

Table 8. *Scaling the hybrid design (sized 3 mm / 1.2 m exchanger) by repeating the 270 kW module (discharge-compliant seawater loop). A module is one 270 kW server room with two cooling paths and six Stage 1 cores; 1 MW is about 3.7 modules and 10 MW about 37.*

Quantity	270 kW module	1 MW (×3.7)	10 MW (×37)
Hybrid cooling electricity (with / without coil bypass)	187–217 MWh/yr	0.69–0.80 GWh/yr	6.9–8.0 GWh/yr
Chiller-only cooling electricity	741 MWh/yr	2.7 GWh/yr	27 GWh/yr
Bay water pumped, hybrid	27,400 m ³ /yr	101,000 m ³ /yr	1.0 million m ³ /yr
Bay water pumped, seawater first	1.13 million m ³ /yr	4.2 million m ³ /yr	42 million m ³ /yr
Stage 1 plate cores (1.2 × 1.2 × 2.7 m)	6	~22	~220
Stage 1 exchanger volume / deck area	23 m ³ / 8.6 m ²	86 m ³ / 32 m ²	860 m ³ / 320 m ²
Cooling electricity cost saved vs chiller only	~\$65,000/yr	~\$240,000/yr	~\$2.4 million/yr

Linear scaling holds only while each module behaves like the one modelled. Four effects break it at facility scale and are not captured here:

- Deck space. Exchangers, filters and ducts need several hundred square metres at 10 MW, and pier deck is costly. Compact corrugated exchangers, which trade fan energy for space, may become the better choice at scale.
- Exhaust recirculation. With many modules side by side, warm exhaust from one unit can reach a neighbour's intake, so interior modules see warmer air than the edge module modelled here. Intake and exhaust placement becomes a site-layout problem.
- Duct length. A centralised plant with long duct runs has higher pressure drop than repeated compact modules; keeping each module self-contained keeps the per-kW fan power found here.
- Absolute Bay impact. Permits respond to total withdrawal and heat. At 10 MW the hybrid pumps about 1.0 million m³ a year against about 42 million m³ for seawater first, so the 41-fold reduction matters more, not less, as the facility grows. Above about 0.65 MW either design's design intake flow passes the federal 316(b) threshold of 2 MGD.

Rack density also changes the picture. AI training racks of 100 kW or more require liquid cooling; the breeze and Bay water would then cool a facility liquid loop rather than room air, a different design that this study does not model.

6.4 Regulatory context

California's Thermal Plan caps new discharges to San Francisco Bay at 4 °F (2.22 K) above natural water temperature [27], the state restricts once-through cooling with coastal water [31], and the Alameda rejection shows the local sensitivity. Section 4.7 shows that a variable-speed pump can hold the modelled coil-loop rise within 2.0 K; actual compliance would require project-specific intake, discharge-plume and permitting analysis. The hybrid pumps in about 15% of hours (about 4% at 24 °C) and returns about 41 times less heat and water than a compliant seawater-first design, which is likely to ease environmental review. That remains a hypothesis until entrainment, mixing-zone and plume conditions are examined with the San Francisco Bay Regional Water Quality Control Board. The conclusion is specific to enclosed waters with a limit at the point of discharge. On open coasts, where limits may apply after mixing and outlets can be tempered with extra seawater as at Hamina [9], seawater-first cooling may remain the better choice. Thermal discharges also affect coastal waters more broadly, for example by reducing CO₂ uptake [32], which is part of why such limits exist.

6.5 Other coastal sites

Industry is moving toward outdoor-air heat rejection at scale: in October 2026 LG agreed to supply air-cooled chillers with integrated free cooling for more than 5 GW of North American AI data-center capacity [33]. The method is not tied to the Golden Gate. It needs three site inputs: hourly air temperature, hourly water temperature and the local discharge rule. North of the Bay, coastal headlands such as Point Reyes and Point Arena lie in the summer upwelling zone of the California Current, where both the air and the near-shore water are typically cooler than inside the Bay, so the air stage would carry a larger share of hours. These are exposed open-coast sites, however: wave loading, access, grid and fibre connections, and marine-sanctuary permitting all become harder, and, as Section 6.4 notes, an open-coast discharge rule may favour seawater first. The same reasoning extends to other cold-current coasts, such as the Pacific Northwest, Chile and the North Sea. For any such site, the first step is to re-run the open-source model with local air, water and discharge inputs.

6.6 Limitations

- Duct, filter and coil pressure drops and fan efficiency are typical design values, not vendor data. The Stage 1 exchanger is sized from laminar theory, which ignores inlet effects, plate spacers, manifold losses and fouling; a 3 mm outdoor-side gap may need a fouling allowance or a wider gap.
- The heat-exchanger model uses the CFD return temperature, which implies about 292 kW of heat removal against a 270 kW IT load; this margin covers the 18–23 kW of room-fan heat (Section 4.7).
- Air data come from one reanalysis grid point and water data from one tide station, with water gaps after February 2024 filled by climatology; three years is short of a climate normal.
- Full cold-aisle containment removes the row-end and over-rack leakage seen in Test 1; the CFD imposes rack airflow and heat rather than resolving server internals. Condensation is avoided only if the room dew point is held at or below about 12 °C.
- The corrosion cost model is illustrative and its floor function is not independently validated.

Capital cost is outside the scope of this paper. Both designs need the same pier, splash-zone protection and server room, so the comparison that matters is the incremental cost: the hybrid adds Stage 1 cores and outdoor-air ducting, while seawater first needs a somewhat larger seawater loop (peak about 24 against 18 kg/s per path under the discharge limit) that runs every hour rather than about 15% of hours, which drives its maintenance and pump wear (Section 5). A payback estimate needs vendor prices for the cores and the seawater loop.

7. Conclusion and Next Steps

For a 270 kW server room on a San Francisco Bay pier, the marine-layer breeze can hold a 21 °C supply in about 85% of hours, with seawater covering nearly all the rest and a chiller needed in fewer than 15 hours a year. The hybrid uses about two-thirds less cooling electricity than a chiller plant. Once the seawater loop is designed to meet California's 2.2 K discharge limit, the hybrid with a sized plate exchanger (67 Pa) and coil bypass uses about the same energy as a seawater-first design (mechanical PUE 1.08 for both) while pumping about 41 times less Bay water. The cost is space, with exchanger stacks about eight times larger than first drawn, an exchanger effectiveness of at least 0.75, which a commercial counterflow core exceeds (81–85% in the maker's selection tool), and well-separated intakes and exhausts. Room CFD confirms that full cold-aisle containment keeps every rack inlet within ASHRAE limits at both 21 °C and 24 °C supply, with results independent of mesh refinement.

Next steps, in order:

1. Obtain vendor-confirmed selections and prices for the Stage 1 cores and the seawater coil, and confirm discharge and intake conditions with the San Francisco Bay Regional Water Quality Control Board.
2. Model a multi-module layout (for example 1 MW) in site-scale CFD to find the actual intake temperature rise from neighbouring exhausts; each 1 K adds 7–11 points of pump runtime.
3. Incorporate review from researchers in data center thermal management and coastal engineering, and target a conference such as ASME InterPACK or IEEE ITherm.

Data and Code

All data, code and figures are available at github.com/sharanankathi/bay-breeze-cooling (MIT licence). The version used in this paper is archived on Zenodo as release v1.0.0, doi:10.5281/zenodo.23204309.

Use of AI Tools

Generative AI (Claude, Anthropic) was used as an assistant to write and debug the analysis code, digitise the CFD line plots, search the literature and edit the text. The author defined the study, built and ran the CFD model, checked all results against hand calculations and source data, and takes full responsibility for the content. AI is not an author of this work.

References

- [1] J. Park, "Designing a very efficient data center," Engineering at Meta, Apr. 14, 2011. [Online]. Available: <https://engineering.fb.com/2011/04/14/core-infra/designing-a-very-efficient-data-center/>
- [2] D. Donnellan, A. Lawrence, D. Bizo, P. Judge, J. O'Brien, J. Davis, M. Smolaks, J. Williams-George and R. Weinschenk, "Uptime Institute Global Data Center Survey 2024," Uptime Institute, 2024.
- [3] D. Swinhoe, "Nautilus Data Technologies launches first floating data center," DatacenterDynamics, Apr. 21, 2021.
- [4] S. Moss, "Nautilus puts Stockton floating barge data center up for sale at \$45m," DatacenterDynamics, Nov. 4, 2024.
- [5] R. Bangert, "Environmentalists sink Nautilus data center," Alameda Point Environmental Report, Jun. 29, 2019.
- [6] S. Tavares, "Alameda rejects controversial waterfront server farm," East Bay Citizen, Jun. 20, 2019.
- [7] DatacenterDynamics, "Start Campus to use Nautilus' cooling tech at Sines Campus in Portugal." [Online]. Available: <https://www.datacenterdynamics.com/en/news/start-campus-to-use-nautilus-cooling-tech-at-sines-campus-in-portugal/>
- [8] I. Latif, M. M. Ashraf, U. Haider, G. Reeves, A. Untaroiu, F. Coelho and D. Browne, "Advancing sustainability in data centers: evaluation of hybrid air/liquid cooling schemes for IT payload using sea water," IEEE, Dec. 2024. [Online]. Available: <https://hdl.handle.net/10919/124238>
- [9] Trellis, "Google uses sea water to cool Finland data center." [Online]. Available: <https://trellis.net/article/google-uses-sea-water-cool-finland-data-center/>
- [10] R. Mokhtari and A. Arabkoohsar, "Feasibility study and multi-objective optimization of seawater cooling systems for data centers: a case study of Caspian Sea," Sustainable Energy Technologies and Assessments, 2021.
- [11] N. L. Boafu, S. K. Anka, S. Boahen and J. M. Choi, "Comparative analysis of the energy-saving potential of a hybrid deep dam water cooling system for data centers," Energy and Buildings, 2025.
- [12] C. S. Chin, J. Zhang and M. Venkateshkumar, "Computing at sea: floating and offshore data centres as a pathway to sustainable AI infrastructure," arXiv:2609.12511, 2026.
- [13] DatacenterDynamics, "HiCloud's offshore wind-powered underwater data center up and running off coast of Shanghai, China." [Online]. Available: <https://www.datacenterdynamics.com/en/news/hiclouds-offshore-wind-powered-underwater-data-center-up-and-running-off-coast-of-shanghai-china/>

- [14] G. Alston, N. Love and R. Haider, "Data center cooling choices shift water impacts across the grid: an integrated water-energy model for sustainable data center development," arXiv:2609.25437, 2026.
- [15] H. R. Hanson and D. C. Hurst, "Corrosion control—offshore platforms," in Proc. Offshore Technology Conference, Houston, TX, 1969, paper OTC-1042-MS. (Also published in JOM, vol. 22, no. 4, pp. 46–50, 1970.)
- [16] C. A. Powell and H. T. Michels, "Copper-nickel: splash zone sheathing of oil-gas platforms," Copper Development Association. [Online]. Available: https://www.copper.org/applications/marine/cuni/applications/offshore_oil_and_gas/splash_zone_sheathing.html
- [17] DNV, "Cathodic protection design," Recommended Practice DNV-RP-B401, May 2021.
- [18] F. Ascione, A. Leonardi, G. Maselli and A. Nesticò, "Advancing sustainable offshore structures through composites materials: structural and economic analysis," Composites Part A: Applied Science and Manufacturing, vol. 209, art. 110008, 2026, doi: 10.1016/j.compositesa.2026.110008.
- [19] J. Clidaras and D. W. Stiver, "Water-based data center," U.S. Patent 8,853,872 B2, Oct. 7, 2014.
- [20] K. Lindberg, R. Longwell, D. Mik and D. Reynolds, "Heat exchanger having longitudinal structure and mounting for placement in seawater under piers for heating and cooling of buildings," U.S. Patent 6,736,191 B1, May 18, 2004.
- [21] P. Zippenfenig, "Open-Meteo.com Weather API," Zenodo, doi: 10.5281/zenodo.7970649.
- [22] H. Hersbach et al., "The ERA5 global reanalysis," Quarterly Journal of the Royal Meteorological Society, vol. 146, no. 730, pp. 1999–2049, 2020, doi: 10.1002/qj.3803.
- [23] NOAA Center for Operational Oceanographic Products and Services, "Station 9414290, San Francisco, CA: water temperature," Tides and Currents. Data accessed 2026.
- [24] ASHRAE Technical Committee 9.9, Thermal Guidelines for Data Processing Environments, 5th ed. Atlanta, GA: ASHRAE, 2021.
- [25] T. L. Bergman, A. S. Lavine, F. P. Incropera and D. P. DeWitt, Fundamentals of Heat and Mass Transfer, 7th ed. Hoboken, NJ: Wiley, 2011.
- [26] R. K. Shah and A. L. London, Laminar Flow Forced Convection in Ducts. New York: Academic Press, 1978.
- [27] State Water Resources Control Board, Water Quality Control Plan for Control of Temperature in the Coastal and Interstate Waters and Enclosed Bays and Estuaries of California (Thermal Plan), 1972, as amended.
- [28] RECUTECH s.r.o., "rSelect online calculator, REC+120 aluminium counterflow heat exchanger," recutechcalculator.com, accessed Oct. 3, 2026.
- [29] Code of Federal Regulations, 40 CFR 125.81 (applicability to new facilities with design intake flow greater than 2 MGD) and 40 CFR 125.84 (through-screen velocity).
- [30] Q. Wei, J. Lu, X. Xia, B. Zhang, X. Ying and L. Li, "Performance and applicability analysis of indirect evaporative cooling units in data centers across various humidity regions," Buildings, vol. 14, no. 11, art. 3623, 2024.
- [31] California State Water Resources Control Board, Water Quality Control Policy on the Use of Coastal and Estuarine Waters for Power Plant Cooling, adopted 2010, as amended.
- [32] Y. H. Ko and T.-W. Kim, "Assessment of the potential effect of thermal effluents on CO₂ absorption in coastal waters," Frontiers in Marine Science, vol. 11, 2024, doi: 10.3389/fmars.2024.1338832.
- [33] LG Electronics, "LG Electronics secures supply agreement to advance AI data center cooling business in North America," Oct. 5, 2026. [Online]. Available: <https://www.lg.com/global/newsroom/news/eco-solution/lg-electronics-secures-supply-agreement-to-advance-ai-data-center-cooling-business-in-north-america/>