

Why AI compute should go to sea

Niccolo Cymbalist, PhD; Kai Matsuka, PhD; Luca Cymbalist

September 2026

Summary

The demand for AI computing capacity is growing exponentially; however, limitations on the availability of power and opposition from local populations are hampering the construction of new terrestrial data centers. This is pushing data center development into alternative, hypothetical environments, including space. We believe that the ocean is another option. It is rich in available energy, immediately accessible, offers effectively unlimited cooling (seawater), and involves more predictable and streamlined permitting and regulations than terrestrial data centers.

We propose a class of vessel that harvests energy from the open ocean at capacity factors approaching 90%. It will use a high-lift wind propulsion system coupled with a hydrokinetic turbine system and weather prediction data to route the vessel along a high-energy path. Energy harvested from the ocean will be used to power onboard AI computing and heat-exchanged seawater will be used for cooling. The construction and deployment of a fleet of these vessels will unlock gigawatt-scale AI infrastructure and power its buildout with clean and effectively unlimited ocean energy. Clippership has identified immediately available shipyard annual production capacity for 300 MW/year worth of on-ocean data center vessels, with expansion potential to 1GW/year. The resulting levelized cost of energy (LCOE) of the power generation system is approximately \$28-35/MWh.

The problem

The world's demand for AI compute capacity is surging, while our ability to build land-based data center infrastructure to meet this demand is limited by deliverable power (generation, distribution, and interconnect) and by widespread opposition to new terrestrial data center construction, particularly in the United States. These issues will become more pronounced as the scale of data center build-out accelerates.

US data center power demand is predicted to increase from 31 GW in 2025 to 66 GW in 2027 [[Goldman Sachs](#)], to 118 GW in 2030 and 194 GW in 2035 [[Bloomberg NEF](#)]. AI agent inference token usage is expected to hit 120 quadrillion monthly by 2030, a 24x increase versus 2026 baseline [[Goldman Sachs](#)].

Power availability is constrained by generation interconnect, load-to-grid connection, and generation & distribution equipment availability timelines. New generation typically takes more than 5 years from interconnection request to commercial operation, and some utilities quote multi-year waits for large-load connections (up to 7 years for >100 MW data centers at Dominion). DOE reports distribution-transformer lead times of one to two years or longer, while some gas-turbine orders have faced substantially longer waits. Advanced nuclear projects remain primarily in licensing or early

construction, so the timing and cost of repeatable behind-the-meter deployments remain uncertain. [[Berkeley Lab](#), [DOE](#), [NRC](#)]

Data center construction pushback at the municipal and state level has rapidly increased in 2026, with data center construction bans and moratoriums increasing to over 350 as of September 2026, including the statewide moratorium in New York and Texas's freeze on new grid-connected projects. It is likely this pushback will intensify as the scale of AI data center buildout accelerates.

These constraints are driving interest in alternatives to terrestrial data centers. The ocean offers plentiful cooling water, tens of terawatts of accessible wind and wave power [[IEA](#)], and access without launch vehicles. Offshore infrastructure occupies sparsely used areas and may face less local opposition far from coastlines. Increasingly available gigabit satellite links provide connectivity.

The Clippership solution

Modern large sailing vessels harvest megawatts of power from the wind, using it to propel the vessel through the water, loaded with cargo or passengers. Most of this power can be recovered and converted to electricity using hydrokinetic water turbines.

Wind-powered ships driving underwater turbines have been proposed since the 1980s, and tested as recently as 2024 by Mitsui OSK Lines, the large Japanese shipping line [[MOL](#)]. This concept suffers from a severe 'stranded power' issue – or how to transfer the generated power to consumers. Hydrogen has frequently been proposed as the medium. This requires complex onboard hydrogen production and storage equipment and incurs efficiency losses approaching of 70% [[CPUC](#)]. Using electrical power directly onboard to power AI compute eliminates the stranded power issue, and unlocks abundant clean ocean power adjacent to virtually unlimited cooling water.



Figure 1. Conceptual 3MWe vessel diagram. 70m LOA, 4x250sqm wing sails, 2x5.5m hydrokinetic turbines (red).

Clippership is developing autonomous wind-powered vessels with hydrokinetic turbines rated at 3 MWe¹, supporting 1,650 late-generation (1200W TDP) accelerators² (see Figure 1). Weather-based

¹ Throughout this paper, MWe references electrical power generation or consumption, while 'MW compute' references cumulative accelerator thermal design power (TDP).

² An NVIDIA GB200 accelerator rated at 1200 W TDP, for example, typically consumes about 85–95% of TDP during heavy, consistent token generation. Non-GPU components add an additional ~4.5 kW per node of 8

routing yields simulated capacity factors approaching 90%, versus the cited 50% for offshore wind [IEA]. Conventional marine LNG generation provides backup power. Table 1 reports capacity factor and intermittency from four years of routing simulations (2022–2025) in an area between California, Hawaii, and Alaska, using NOAA GFS 0.25-degree historical forecast data [NCEP], using a simple 7-day look-ahead routing algorithm. Visual simulation and routing output over one year is available [here](#). P75 and P90 are the power levels exceeded 75% and 90% of the time, respectively, in each simulated year.

Table 1. Capacity factor and intermittency metrics for 3MW nameplate vessel operating in the Pacific Ocean

Year	Cap. factor	Hours > 2.7 MWe	Hours < 2 MWe	P75 power, MWe	P90 power, MWe
2022	87.92%	6,234	1,362	2.592	1.563
2023	89.25%	6,570	1,110	2.700	1.691
2024	88.37%	6,258	1,290	2.605	1.674
2025	88.12%	6,168	1,260	2.597	1.660




Figure 2. Year-long (2025-2026) simulated Pacific route achieving capacity factor of ~90% vs nameplate 3MW hydrokinetic power recovery [full route simulation]

The vessels will be designed to accept standard modular data center compute units in 40’ containers (Crusoe Spark or Armada Galleon units, for example).

Clippership estimates 160,000 gross tons of vessels per gigawatt of generation capacity, roughly one large cruise ship. A 10 GWe fleet would therefore total approximately 1.6 million GT. For comparison,

GPUs. The sum is then multiplied by the facility PUE (1.05 in our case, thanks to seawater cooling), resulting in an electrical peak power consumption of 1750 W. [Spheron]

South Korean shipbuilding capacity is about 18 million GT annually [OECD]. These figures provide a benchmark for the scale of construction required; identified shipyard capacity is detailed below.

Vessel and system architecture

At its core, the Clippership vessel is a system that uses the thrust from the wind to overcome both hull resistance and induced hydrokinetic turbine resistance, or:

$$F_{drive} = R_h + R_t$$

F_{drive} is the wing thrust, R_h is hull and appendage resistance, and R_t is turbine resistance. A simplified interactive power calculation model can be found [here](#).

Based on this model, for a conceptual 70-meter, 3MWe-class catamaran vessel with 1000 square meters of active suction wing ($4 \times 250 \text{ m}^2$, $C_L = 7$) in 10.5 m/s of true wind on a beam reach and traveling at 16 knots, a representative force balance then becomes:

$$570 \text{ kN } (F_{drive}) = 100 \text{ kN } (R_h) + 470 \text{ kN } (R_t)$$

To harness the power of the wind, convert it to usable electrical power, and house the powered payload, Clippership vessels have three main subsystems:

1. **The rigid wing system** that extracts power from the wind and converts it into vessel forward thrust. A medium-sized 70-meter (3MWe rated power generation capacity) vessel will have about 1000 square meters of wing area. Clippership has developed and tested this wing system over the past two years, and it is now undergoing extensive open-ocean testing. Current development efforts are focused on increasing the lift coefficient of the passive wing from the current $CL = 2.6$ to $CL \approx 7$ using active suction systems comparable to those developed by bound4blue (stated maximum $CL = 8.3$) and others.




Figure 3. Clippership wing system and autonomous vessel hull undergoing wind tunnel and on-ocean testing in strong wind conditions.

2. **The hydrokinetic generation system** that extracts power from the movement of the vessel through the water. This consists of an underwater 7-meter diameter axial hydrokinetic turbine

driving a 3MWe generator at a vessel speed of 8 m/s (or two 5.5-meter diameter 1.5MWe turbines). The hydrokinetic turbine diameter is considerably smaller than that of a wind turbine of equivalent power (7-meter versus 130-meter diameter) due mainly to the 800-fold density difference between air and water. Hydrokinetic turbine components do not require specialized supply chain and are based on existing hydrokinetic turbines developed for tidal power projects, for example those deployed by [Orbital Marine](#), that share considerable overlap with wind turbine supply chains.

3. **The autonomous hull**, containing the data hall, power electronics, heat-exchanged cooling loops, backup power generation equipment, autonomous control systems, communication and navigation equipment. Notably, given abundant cool (15–18 °C) seawater available in the planned area of operation for heat-exchanged liquid cooling, anticipated PUE is on the order of 1.05. The hulls are designed to be manufactured at shipyards anywhere in the US or abroad. We have identified shipyard capacity to achieve 300 MWe annual production.

The economics

A useful approach to comparing the costs of the proposed on-ocean compute vessels to terrestrial data centers is to split out the power generation system (including wing systems, hydrokinetic generator, and associated power electronics) from the vessel itself, which houses the compute hardware, cooling systems, backup power generation, UPS, and other systems with direct terrestrial data center analogues. The power generation system can then be assigned its own LCOE, and the rest of the vessel can be benchmarked against terrestrial powered-and-cooled shell costs.

Table 2. 3MW vessel hull construction costs (powered and cooled shell)

Cost Item	Cost	Notes
Hull	\$6,250,000	Based on a hull structure mass of 240 tons, and \$26/kg aluminum construction [quoted].
Outfit excluding backup power generation	\$5,800,000	Includes standard marine outfitting, auxiliary propulsion, electrical system, data hall shell. Excludes IT-specific electrical system.
3 MW power backup system	\$4,700,000	Includes LNG gensets and cryogenic LNG storage tanks, regasification train, and other systems
Project costs & contingency	\$2,000,000	
Total	\$18,750,000	In line with similar (but more complex) vessels [Austal]

Resulting powered and cooled shell cost is on the order of \$6.25M/MWe electrical power capacity. Notably, this is close to an equivalent terrestrial powered and cooled shell per MW electrical power, estimated at between \$5.5M/MWe and \$6.5M/MWe [[Turner & Townsend](#)].

Table 3. 3MW hydrokinetic turbine and wing system (power generation) costs

Cost Item	Cost	Notes
-----------	------	-------

Wing system	\$2,680,000	Modern WASP system designed for manufacturing at scale ³
Hydrokinetic generator system	\$2,550,000	Marinized equivalent wind turbine costing exercise ⁴
Total	\$5,230,000	

Resulting LCOE for the wing + hydrokinetic turbine system is therefore approximately \$28/MWh at the 88% capacity factor achieved in simulation.⁵ Reducing capacity factor to 80% drives up LCOE to \$31/MWh, and 70% to \$35/MWh. This is lower than US onshore and offshore wind (currently \$37-99 and \$105-167 per MWh, respectively [[Lazard](#)]) and is comparable to Chinese onshore wind power at \$29/MWh, one of the lowest-cost energy sources on the planet [[BloombergNEF](#)].

Operating costs for the 70-m autonomous vessel are estimated to be between \$0.8M and \$1.7M per year, based on operating costs of equivalent conventional vessels [[Tidewater](#)] excluding crew costs (\$1.0M–\$1.3M per year), which autonomous operation eliminates. LNG fuel costs needed to achieve 100% capacity factor at full 3MWe load are \$280k/year as of September 2026, and track LNG bunker benchmarks [[S&P Global](#)].

For reference, current base-case terrestrial powered shell revenue is \$2.1M/MW-yr (\$175/kW-month, and increasing every quarter as supply tightens and demand increases, [CBRE](#)) exclusive of electrical power. A single 3MWe vessel can then generate \$6.3M/year as a powered shell today. Selling power at a competitive rate of \$60/MWh increases revenue by \$1.5M.

Table 4. Illustrative payback period sensitivity. Compute hardware is customer owned.

	Poor	Base	Good
Revenue (\$M)	\$6.30	\$7.80	\$11.50
Vessel cost (\$M)	\$28	\$24	\$20
Opex (\$M)	\$2	\$1.50	\$1
Payback (years)	6.5	3.8	1.9

Table 4 shows sensitivity of unit economics to poor revenue conditions (\$140/kW powered shell rental per month, \$50/MWh power), and higher-than-expected vessel cost and Opex, resulting in a

³ The composite wing, flap, and primary structural components account for roughly **\$290k per wing**, while the active-suction system adds approximately **\$130k**. Rotation, actuation, controls, electrical systems, and marine hardware contribute approximately **\$90k**. Assembly, testing, factory overhead, warranty, and manufacturer margin make up the remaining approximately **\$160k per wing**, for a total of **\$2.68M** for the 4x250sqm wings.

⁴ Major mechanical components include the blades (\$360k), hub and pitch system (\$160k), shaft and bearings (\$200k), and gearbox (\$260k), totaling approximately **\$980k**. The generator, power converter, transformer, and associated electrical equipment contribute approximately **\$580k**. Marine-specific systems—including the pressure housing, shaft seals, cooling, controls, cabling, and corrosion protection—add approximately **\$670k**. Final assembly, factory testing, integration, and margin add approximately **\$320k**, for a total of \$2.55M.

⁵ Assumes an initial capital cost of \$5.23 million and a nameplate electrical output of 3 MWe. At a 90% capacity factor, the system would generate approx. 23,652 MWh/yr. Assuming a 25-year economic life and an 8% real WACC, the corresponding capital recovery factor is 9.37% per year, resulting in an annualized capital cost of approximately \$490,000. Fixed operations and maintenance costs are assumed to equal 3% of initial CAPEX, or approximately \$157,000 per year.

payback period of just over 6 years. Conversely, good revenue conditions (\$250/kW monthly powered shell rental, \$100/MWh power offtake), vessel build cost reduction, and a reduction in Opex result in a shorter payback period of about 2 years.

Roadmap

October 2026. Launch our first two dedicated sense and compute vessels, each rated at 10 kWe and powering eight NVIDIA H200 GPUs with heat-exchanged seawater cooling. They will serve maritime and government customers requiring local modeling and analysis, including bathymetry, video and radar processing from hull-mounted sensors or daughter UUVs/USVs, and weather prediction beyond the compute capacity of smaller craft. These vessels validate end-to-end power generation through useful AI workloads and provide early indicative information on uptime and bandwidth.



Figure 4. Sense and compute node. Second hull under construction at partner shipyard (right).

Q2 2027. Launch a scaled-up 0.25 MWe-class data center vessel. It will be used as a platform to test and validate the scaled-up power generation stack (rigid wing system through hydrokinetic turbine), and the routing optimization, with an emphasis on demonstrating the capacity factor and uptime over a representative period, in a representative area of operations. This vessel will be able to host, power, and cool approximately 150 accelerators, and will be available for limited commercial deployment for customers requiring isolated, sovereign AI compute resources.

Q2 2028. Launch the first 3 MWe-class prototype. This class of vessel will have the capability to act as a drop-in replacement for latency-insensitive inference workloads currently assigned to terrestrial data centers. It will validate the full-scale node end to end.

Q4 2029. Begin production of the 3 MWe-class data center vessel once the design has been validated and frozen, targeting 1 GWe (~330 vessels) per year in the early 2030s.

The hard questions

How will you handle bandwidth and latency limitations? How will they affect serviceable AI workload?

SpaceX has filed plans with the FCC to launch up to 100,000 3rd generation Starlink satellites into VLEO, offering multi-gigabit symmetrical throughput. Amazon is targeting gigabit speeds with its Amazon Leo constellation [[FCC](#), [Amazon](#)]. Blue Origin plans to roll out its TeraWave product capable of symmetrical 144 Gbps in RF starting in Q4 2027 [[Blue Origin](#)].

The 0.25 MWe and 3 MWe prototype vessels will use the currently available 6x Bonded Performance Starlink service achieving gigabit speeds (2.4 Gbps downlink and 600 Mbps uplink), with latency between 50 and 100 ms. Vessels requiring higher bandwidth earlier than next-generation gigabit options become widely available will incorporate Starlink Gateway units that are specifically designed for telecom infrastructure (including data center) connectivity, and have been successfully deployed in maritime environments [[SpaceX](#)].

Recent inference studies (based on tok/s and GPU power) report between 1 and 2 Mbps/kW compute ingress bandwidth requirements, and about 100-fold lower egress bandwidth requirements for agentic workloads [[DeepSeek](#), [SemiAnalysis](#)]. For a 3MWe vessel powering approximately 2 MW of compute hardware, this would require between 2 and 4 Gbps ingress, and sub-500 Mbps egress.

How will you ensure the onboard hardware remains safe and protected from the marine environment? How is it serviced?

Onboard IT hardware will be housed in a sealed and nitrogen-purged data hall. Microsoft's Project Natick demonstrated the effectiveness of a nitrogen-purged sealed environment in a marine data center pilot, achieving a failure rate one-eighth that of terrestrial data centers. The vessels will periodically navigate, under backup power, to service ports located on the closest suitable coast. There, vessel technicians will perform routine vessel maintenance tasks, while concurrent data center hardware maintenance and updates can be performed.

How is this better than offshore wind turbines, advanced wave energy solutions, or offshore nuclear?

The ability of the vessels to move under their own power unlocks capacity factors approaching 90%, almost double those currently achieved by floating or fixed turbine-based offshore wind installations. It also effectively eliminates any offshore construction activity that is a major cost driver of offshore wind installations. The vessels operate in international deep water away from coastlines and are therefore not subject to the same public scrutiny and political pushback faced by offshore wind installations.

Advanced wave energy also benefits from the advantages of international deep-water operation, but limited mobility of the platform reduces capacity factor and complicates both outbound transit to the area of operation and return transit for maintenance and overhaul.

Offshore nuclear is attractive, but is reliant on widespread deployment of marine SMR technology whose development has only just started, whose build and operating costs are still unclear, and that is unlikely to see any substantial deployment in the next decade. The vessels we are developing leverage existing technology and marine concepts of operations (CONOPS), enhanced by the wing-sail energy harvesting systems that we have already developed and demonstrated.

What is the regulatory environment and roadmap to deploying gigawatt-scale on-ocean compute?

The regulatory pathway for our autonomous surface vehicles is navigable, and actively being built out by regulators. Our vessels route through the dedicated review track the USCG created for uncrewed and autonomous systems in its Work Instruction, published in June 2026 [[USCG CG-5P](#)]. We have been in sustained engagement with U.S. and international maritime-autonomy regulators since 2024 and are executing a phased deployment plan for uncrewed autonomous vessels. We engage closely with USCG Sector San Francisco for our routine autonomous operations in the Bay and offshore. [[USCG Clippership LNM](#)]

How fast can you deploy? 1 GWe per year corresponds to ~330 vessels per year, which is a lot!

We have identified and are in contact with 300 MWe/year of shipyard capacity available after design validation and freeze across six yards. In some cases, these yards are already building our vessels and report capacity of up to twenty 3 MW vessels each annually. Our planning total is approximately 100 vessels per year, with typical build times of 12–14 months. We have identified additional yards capable of expanding to approximately 1GW/year worth of capacity.

The wing system and hydrokinetic turbine system share a supply chain with utility-scale wind turbines (blades and nacelles), whose installed manufacturing capacity approaches 285 GW/year.

Conclusion

The large and readily exploitable sources of energy and cooling provided by the ocean present an excellent opportunity for developing ocean-based AI compute infrastructure that is strongly competitive with terrestrial data centers. Add the relative regulatory freedom of the open ocean compared to the complex challenges faced by terrestrial data centers, and ocean-based AI compute appears to be the logical next step in AI infrastructure buildout.

List of interactive visualizations

1. [One-year route simulation](#)
2. [Vessel power model](#)

About Clippership

We are a team of sailors, aerodynamicists, mechanical, electrical, and software engineers and technicians, and SOF veterans dedicated to harnessing the power of the oceans to help solve the problems humanity faces today. In our previous lives, we built humanoid robots and rocket-powered cars at Tesla, autonomous spacecraft at NASA-JPL, race cars with the Mercedes F1 team, ships at Damen, codebases at Microsoft, and have deployed in combat theaters around the world. We're proud to build in South San Francisco. Contact us at hello@clippership.co.