HYDRO WIND ENERGY / REPORT 2026

Infrastructure for the AI era

The Offshore AI Power Plant

Renewable power, long-duration mechanical storage, deep-ocean cooling, and floating data-centre infrastructure — integrated into a single offshore platform.

hw.energy

CONTENTS

Inside this report

01 / EXECUTIVE SUMMARY

AI is no longer just a software revolution. It is an infrastructure race.

The next constraint for AI is not chips or models — it is the physical systems required to power, cool, scale, and operate next-generation data centres.

Hydro Wind Energy is building that layer where energy and cooling are abundant: offshore.

THE CORE IDEA

Move AI infrastructure to where the energy and cooling already are.

02 / THE PROBLEM

AI is now an infrastructure race

Large-scale training, inference, automation, and AI-native enterprise systems all demand reliable power, high-performance cooling, low-cost scalability, and rapid deployment — faster than the current model can support.

> "Traditional cloud infrastructure was not designed for the age of intelligent automation."
>
> Crusoe · 2026 AI Infrastructure Trends

03 / THE PROBLEM

Six constraints on AI infrastructure

01

Power availability

Grid connections at AI scale are delayed, expensive, or simply unavailable.

02

Cooling demand

High-density compute produces heat that consumes energy, space, and water.

03

Land & permitting

Campuses need land, substations, water access, and long permitting timelines.

04

Cost volatility

Unpredictable pricing, hidden fees, and high charges for scaling.

05

Lack of control

98% of decision-makers rate complete control over their data centres as important.

06

Sustainability pressure

Growing demand for transparency on how AI workloads are powered.

EXHIBIT 01 · LIVE FROM THE GRID, JULY 2026

The onshore model is already straining

As a record heat dome settled over the US, PJM — the nation's largest grid operator, serving 13 states — asked that data centres switch to backup generators within 15 minutes of an emergency signal to free up power for homes.

9%	**40%**	**5M**	**7 in 10**
of US power demand from data centres by 2030, up from 4% today.	of a data centre's electricity goes to cooling alone — rising as temperatures climb.	gallons of water a single large AI data centre can consume each day.	Americans oppose new data centres being built in their community.
US Dept. of Energy	Gasilov Group	EESI	Gallup

Source: Al Jazeera, "US heatwave raises alarms over AI data centre energy demands," 3 July 2026

EXHIBIT 02 · MARKET SCALE

The build-out is the largest in history

Capital is committing to the physical stack behind AI — compute, data centres, and power — faster than energy and land can be delivered. The bottleneck is no longer chips; it is infrastructure.

$7.6T	**945TWh**	**~100GW**	**45GW**
of AI-related capex forecast between 2026 and 2031, across compute, data centres, and power.	global data-centre electricity demand by 2030 — more than double the 415 TWh of 2024.	of new data-centre capacity added 2026–2030 — effectively doubling global capacity.	projected US data-centre power shortfall by 2028 — supply cannot keep pace with demand.
Goldman Sachs	IEA, Energy & AI	JLL Data Center Outlook	Goldman Sachs

In the US, data centres are on course to account for nearly half of all electricity-demand growth to 2030 — concentrated in exactly the regions where grids and water are already stretched.

Sources: IEA, Energy & AI (2025) · Goldman Sachs Global Institute (2026) · JLL 2026 Global Data Center Outlook

04 / THE SOLUTION

One integrated system, built around the needs of AI

Power · Storage · Cooling
Scalability · Control · Sustainability

01

Renewable power generation

OceanHydro Omni captures offshore wind and converts it into mechanical energy.

02

Long-duration storage

Submerged mass-based storage smooths intermittent wind into reliable output.

03

Deep-ocean cooling

Cold seawater from depth cools compute through closed-loop heat exchange.

04

Floating AI data centres

Modular units powered and cooled directly offshore — no land or grid upgrade.

A NEW CATEGORY OF INFRASTRUCTURE

An offshore AI factory — powered by renewable energy, cooled by the ocean.

The emerging "AI factory" is a vertically integrated platform combining energy, data-centre construction, and managed services. Hydro Wind Energy applies that logic offshore — with power and cooling built in from day one.

05 / CAPABILITIES

What the platform does

Generates

Captures offshore wind via OceanHydro Omni.

Stores

Mechanical storage via submerged masses, not only batteries.

Stabilises

Delivers more predictable power blocks to compute.

Cools

Uses deep-ocean water as a natural cold sink.

Unburdens

Cuts dependence on scarce land, grid queues, and freshwater.

Scales

Expands from demonstrator to gigawatt-scale hubs.

06 / RATIONALE

Why offshore

80%

of the world's wind energy is in deep offshore waters

Abundant energy

Offshore wind is stronger and steadier — and HWE builds storage into the generation platform itself.

∞

continuous cold seawater at depth

Natural cooling

Deep seawater is a low-temperature cold sink, cutting the burden of chillers and water-intensive cooling.

GW

scalable to gigawatt-class hubs

Space for scale

Deploy free of urban limits, land availability, and grid congestion that cap onshore campuses.

EXHIBIT 03 · PROOF POINT

The model is already operating

In 2026, the world's first offshore wind-powered, seawater-cooled data centre reached full commercial operation off Shanghai — a 24 MW facility running ~2,000 servers on offshore wind, cooled passively by the ocean, with a second stage planned at 500 MW.

It proves the thesis. Hydro Wind Energy goes further — integrating generation, long-duration storage, and cooling into one serviceable floating platform, rather than a fixed subsea capsule.

<1.15 PUE achieved — vs ~1.5 for typical land-based centres.	**−22.8%** less electricity consumed versus a comparable land facility.
0 freshwater used — cooling draws entirely on seawater.	**−90%** less land required than an equivalent onshore campus.

Source: Xinhua, Data Center Dynamics & Tom's Hardware on the Lingang (Shanghai) offshore underwater data centre, 2025–2026. Comparable floating projects: Panthalassa, Aikido; precedent: Microsoft Project Natick.

07 / ARCHITECTURE

The platform, from surface to seabed

L1	**Renewable Power Capture**	OceanHydro Omni captures offshore wind and converts it into mechanical energy.
L2	**Mechanical Energy Storage**	Submerged masses are raised and lowered to store and release energy on demand.
L3	**Power Management**	Power electronics and controls regulate delivery to the data-centre load.
L4	**Deep-Ocean Cooling**	Closed-loop heat exchange transfers compute heat to cold seawater.
L5	**Floating AI Infrastructure**	Modular data-centre units house compute, networking, and supporting systems.
L6	**Freshwater Production** · optional	Desalination adds strategic value in offshore and remote deployments.

08 / ROADMAP

A staged path to gigawatt scale

P1

Subscale

Offshore validation of performance and survivability.

P2

500 kW–1 MW

Demonstrator proving generation, storage, and delivery.

P3

Integrated pilot

Power, storage, cooling, and compute combined.

P4

10–25 MW

Commercial pilot for a strategic partner or programme.

P5

100 MW

Full commercial platform for large AI workloads.

P6

Gigawatt hubs

Offshore AI zones serving regional and global demand.

09 / VALUE

The value proposition

OPERATORS	CLOUD	GOVERNMENTS	INVESTORS	ESG
Dedicated renewable power, storage, and cooling — no waiting for grid upgrades.	A new route to scalable AI capacity where land and energy are bottlenecks.	Sovereign AI capacity with less pressure on national grids and freshwater.	A new infrastructure asset class across AI, renewables, and ocean tech.	A clear link between AI workloads and renewable energy sourcing.

10 / DIFFERENTIATION

How it is different

vs. land-based data centres	Reduces reliance on land, freshwater, and congested power grids.
vs. conventional offshore wind	Uses energy at the point of generation instead of exporting it to shore.
vs. batteries	Integrates long-duration mechanical storage for energy shifting.
vs. hyperscaler infrastructure	Power, cooling, and control integrated from the ground up.
vs. subsea data centres	A floating platform designed for serviceability, modularity, and expansion.

11 / COMMERCIAL

Commercial models

01

Power & cooling-as-a-service

Dedicated renewable power and cooling under a long-term service agreement.

02

Offshore PPA

A cloud, AI, or data-centre customer contracts for renewable power.

03

Joint venture

HWE provides power and marine infrastructure; the partner runs compute.

04

Sovereign partnership

A national entity backs deployment to create domestic AI capacity.

05

Infrastructure fund

A long-duration asset backed by contracted AI power and cooling revenue.

POSITIONING

The energy-first layer for sovereign, scalable, sustainable AI.

12 / TIMING

Four constraints, solved at once

01 Energy
Renewable power generated offshore.

02 Storage
Long-duration mechanical storage integrated.

03 Cooling
Deep-ocean thermal resources.

04 Land & grid
Infrastructure moved offshore.

THE MESSAGE

The AI infrastructure race will not be won by compute alone — but by those who control the power, cooling, and storage beneath it.

Hydro Wind Energy is building that physical layer offshore.

A NEW CLASS OF INFRASTRUCTURE

The world's first Offshore AI Power Plant — renewable, dispatchable, ocean-cooled, modular.

A vertically integrated renewable power, storage, cooling, and floating data-centre platform — unlocking scalable AI beyond the limits of land, grid, and water.

— Hydro Wind Energy · hw.energy