Subsea Cloud

Subsea Data Centers

Presented by: Maxie Reynolds

SUBSEA CLOUD

ARCHITECTURE, LOAD BALANCING AND REDUNDANCY

How the compute is organized, how load moves when hardware fails, and why availability holds at Tier 4.

How the compute is organized

Redundancy is layered from the server up, so no single failure reaches the service.

- Module (POD): standardized pressure-compensated container. 500 kW to 1.8MW IT in the current configuration, 150+ kW per rack, rated for 10 to 9,000 ft of water.
- Five subsystems per module: compute, internal power (converter, busbar, UPS bridge, protection), passive thermal path, pressure boundary, controls.
- Servers run in active 1-of-2 parallel pairs. A failed node is absorbed in software; nothing gets retrieved.
- Modules share load in active k-of-n parallel. Losing one module trims capacity by 1/n; the service keeps running.

Cluster: modules in active k-of-n parallel

Module A

- Compute (1-of-2)
- Internal power
- Passive thermal
- Pressure boundary
- Controls

Module B

- Compute (1-of-2)
- Internal power
- Passive thermal
- Pressure boundary
- Controls

Shared boundary: 2 diverse subsea cables (6 cores each) + dual-fed topside switchgear

How we load balance

Active-active everywhere. The software layer balances against the capacity tenants use.

- Every module carries live load continuously. There is no idle standby unit on the seabed.
- A node failure is handled by software failover. A module outage, or a planned retrieval, migrates service states to the surviving module without interruption (unless stated otherwise in the SLA).
- Passive cooling keeps junction temperatures low and stable, which lets a module run a continuous ~30% overclock when it has to carry more.
- Reductively speaking: at typical utilization (~60%), a two-module site behaves as if it had a spare module, without paying for one.

$\alpha \approx 1.30$

overclock headroom per module, enabled by passive seawater cooling

650 kW

delivered by a surviving 500 kW module. A 1 MW contract stays whole while active demand is under 650 kW ($U \leq 0.65$); beyond that, excess workloads queue until restoration.

Technical inputs: nothing beyond a standard colo

If a tenant can spec a land deployment, they can spec a subsea one. The marine engineering never crosses their desk.

Tenant provides (same as any provider)

- Capacity and power requirement
- Hardware spec, or their own IT (tenants can supply GPUs in AI/HPC deployments)
- Network and interconnection requirements
- Software stack, hypervisor, workloads

Subsea Cloud handles, end to end

- Engineering and manufacturing
- Site survey, permitting, installation
- Monitoring, security, maintenance
- Power and fiber, seabed to topside

The battery limit sits where a colo's would: tenant owns IT software and workloads (typically), we own everything from the rack to the topside switchgear. Modules take standard server-class equipment. No marinization, no special hardening, no exotic formats.

SUBSEA CLOUD

Nothing beyond standard. What makes our projects easy is what we do not need. No land purchase, no building, no chillers, no water rights, and no fight for grid interconnection in a queue of hyperscalers. We typically apply for inclusion on a partner's already-permitted site, which collapses the permitting timeline and puts the module next to power that is already there. Scaling is the same operation repeated: each additional module(s) is another 12-week build and another few hours down to seabed, tied into the cluster without touching live tenants. Capacity arrives in 500 kW– 1.8 MW increments at roughly $1 million per MW, and costs fall per MW as units deploy.

Redundancy, quantified

99.995%	99.995%	0.48 h/yr
single-module steady-state availability	cluster effective availability, Tier IV	modeled downtime, with topside generator

Compute: active 1-of-2 pair in every module; pair availability 99.9951%

Internal power: UPS bridge with solid-state fault isolation; failover in minutes, not interventions

Cables: 6 cores (min.) per hybrid umbilical, two geographically diverse routes; pair ≈ 99.998%

Topside: dual-fed switchgear, hot-standby spares, 2-day onshore repair; ≈ 99.9996%

Modules: active k-of-n cluster redundancy; overclock headroom covers k = n configurations

A single cable caps any site at 98.90%, no matter what sits below it. Dual corridors are the baseline for that reason.

Method: Weibull life / lognormal repair Monte Carlo, 100,000 ten-year histories. Inputs traced to OREDA, Telcordia SR-332, Quanterion NPRD/EPRD and ICPC. Built to ISO 14224, IEC 61703, IEC 61078 and ISO 20815; common-cause failures per IEC 61508 β-factor.

Traditional vs SSC: maintenance

	Traditional	Subsea Cloud
Cadence	Continuous. On-site staff, rolling preventive and reactive work.	Episodic. Planned campaigns at set threshold, depending on tenant, scheduled on Weibull B10/B50 wear-out thresholds.
Failure response	Dispatch a human to the floor.	Server faults: software failover, no retrieval. Physical faults: intervention, possible within ~12 hours.
Method	In-situ component swaps inside a live environment.	Retrieve, refresh, redeploy. One marine window combines IT refresh, power overhaul and pressure-boundary inspection (unless stated otherwise in SLA).
Monitoring	BMS plus staff walkdowns.	Condition-based telemetry; years 2–4 of the lifecycle are typically remote-only.
Cost shape	Continuous opex, hard to forecast.	Two planned capital events with defined envelopes at Year 5, and Year 10, excluding IT hardware.

Cost envelopes are concept-level allowances from the RAM models, not quotations.

Traditional vs SSC: redundancy

	Traditional	Subsea Cloud
How Tier IV is reached	Duplicate everything: second power train, generators, chillers, idle spares (2N).	Redundancy in the architecture: parallel compute pairs, k-of-n modules, dual cable corridors, dual-fed topside. A module pair behaves like a split busbar.
Cooling	Mechanical plant that can itself fail, so it needs its own redundancy.	Passive seawater convection. No pumps, no chillers, nothing rotating to duplicate.
Spares	Idle hardware waiting for a failure.	Elastic redundancy: ~30% thermal overclock headroom plus utilization slack does the job of an idle spare module.
Capacity sizing	Sized to the worst-case peak of every rack; much of it sits unused.	Software balances against active demand, not sold footprint.
Result	Tier IV at high capex.	The same 99.995%, without the duplicated plant.

Traditional vs SSC: cross connects

	Traditional	Subsea Cloud
Termination	Meet-me room in the building.	Topside fiber cross-connect, inside SSC scope of supply.
Physical path	Single building entry unless diverse entry is bought separately.	Fiber rides the hybrid umbilical alongside power, on two geographically diverse routes, by default.
Provisioning	Ordered per connection, billed as a recurring line item.	Standard interconnection from the topside outward: same carriers, same process a tenant runs anywhere else.
Tenant experience	Standard.	Identical. The difference is buried in the path, not the process.

The cross connect a tenant orders is the one they know. What changes is underneath it: the path to shore is duplicated and physically protected as the baseline design, not an upgrade.

Adding cross connects: module to module

- Interfaces are standardized across modules (mechanical, electrical, software), so any module can join any cluster.
- Module-to-module links run as subsea fiber jumpers, placed by the same ROV tooling we use for maintenance. Hours of marine work, not a construction phase.
- Scaling a site: a new module is built in 12 weeks, fixed to the seabed in as little as an hour once on location, then tied into power and fiber subsea.
- Nothing routes through a live hall. No tenant's traffic is disturbed while capacity is added.

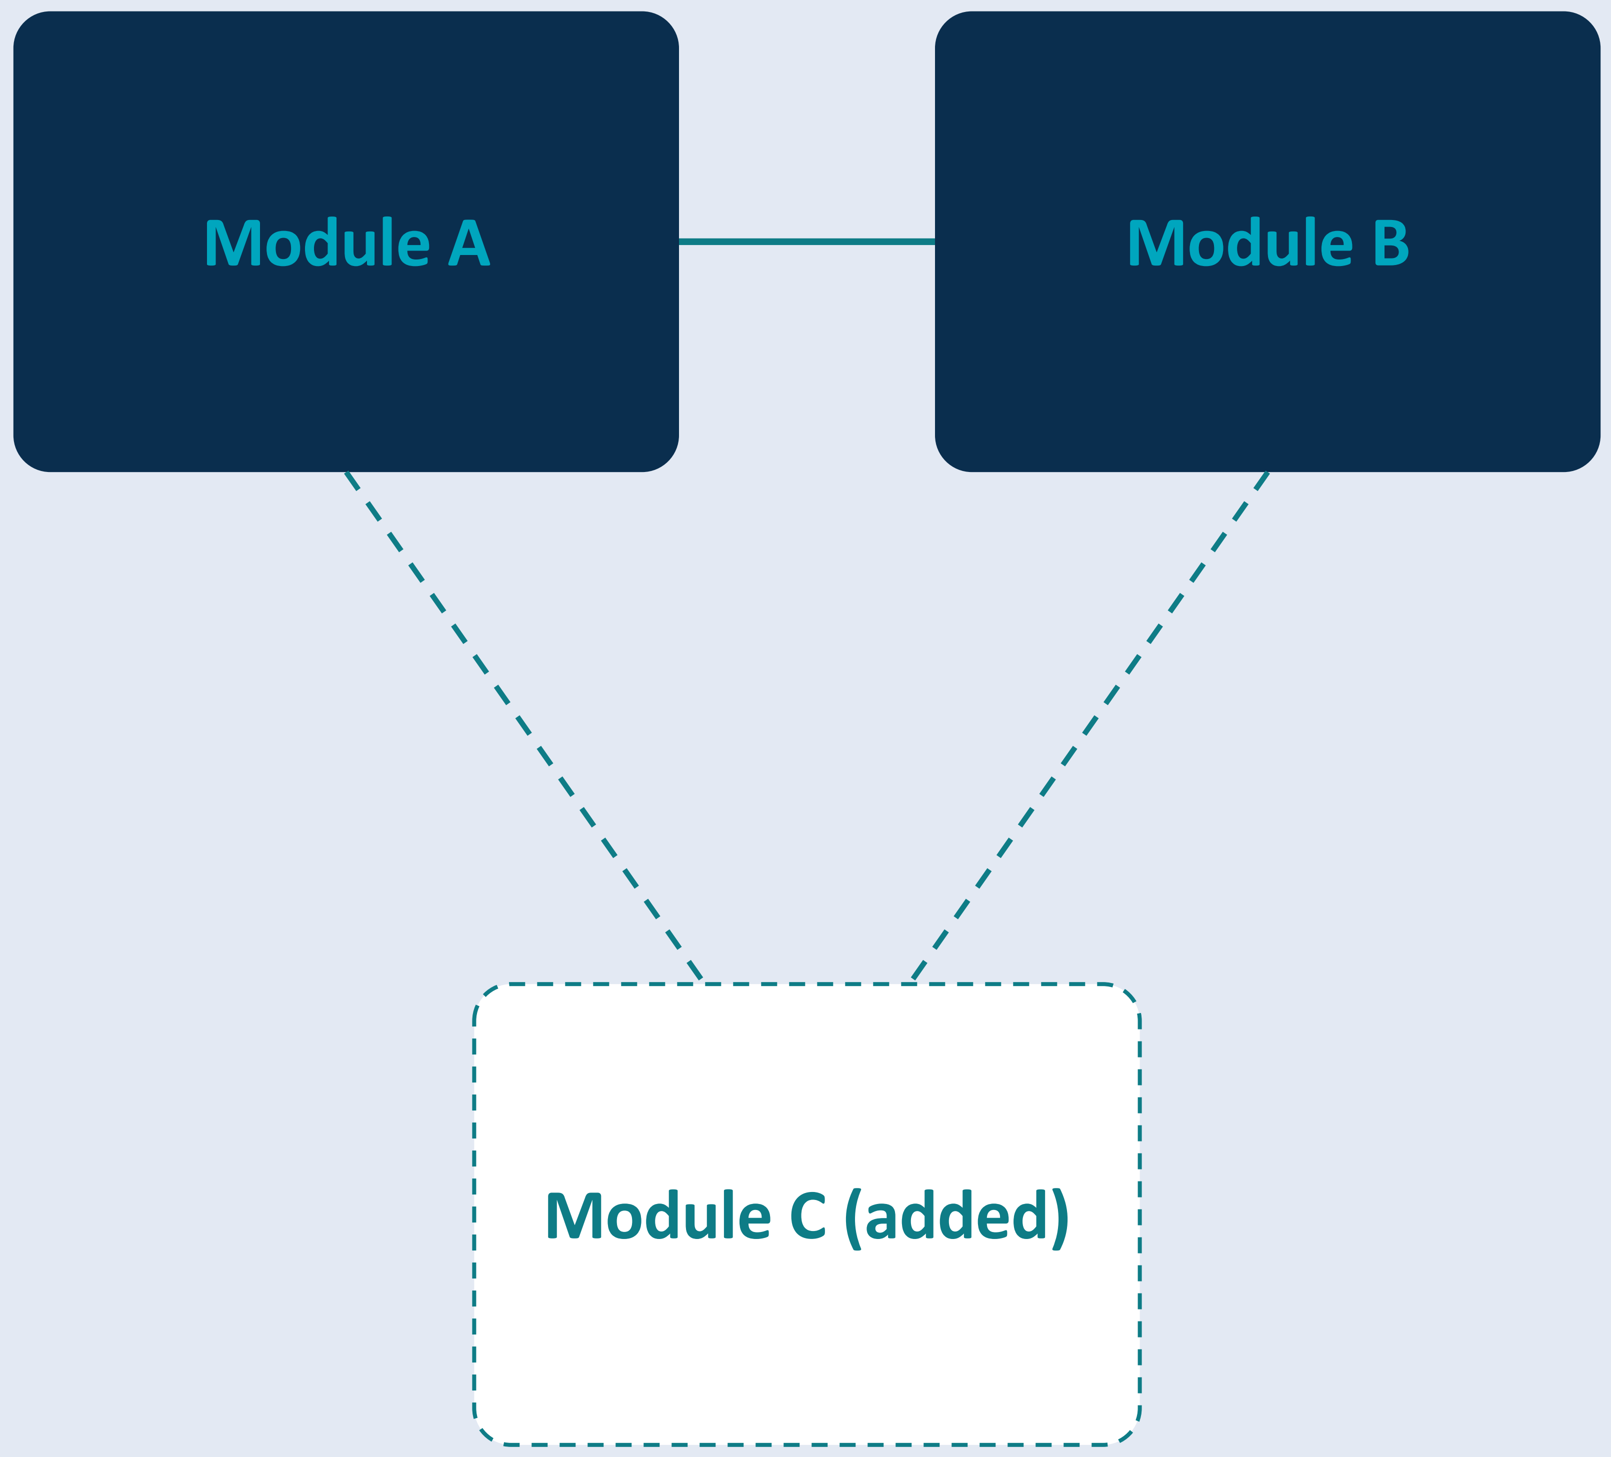

Fiber jumpers, ROV-installed

The rest of the technical case

Electronics last longer cold

Stable junction temperatures. Every 10°C drop roughly halves power-electronics failure rates. We see servers failed at 1/8th the land rate over three years.

Zero water draw

A traditional site uses 3–5 million gallons per day for cooling. We use none, and no refrigerants.

Physical security by default

Sealed, lights-out, on the seabed. No doors, no badge readers, no walk-in attack surface. Used in grey-zone and active defense situations.

Grid outages covered

A topside standby generator holds the power source at 99.9994% and keeps the site at Tier IV on a 99.9% grid.

Site flexibility

10 to 9,000 ft. Rivers, ports, dams, offshore. Co-locates with renewables and existing offshore assets.

Auditable, not asserted

Every availability figure reproduces from public databases: OREDA, Telcordia SR-332, Quanterion, ICPC.

SUBSEA CLOUD

Subsea compute is not a bet on new technology; it is a change of tenant. Oil and gas has been putting electronics on the seabed for more than 50 years: control modules, sensors, power distribution and processing hardware, running at depth, for decades at a time. That history is deep enough that its failure rates fill public engineering databases & OREDA data.

SUBSEA CLOUD

Nearly everything you will touch today is a petrochemical product, from the rubber on your soles to the keys you type on to the electricity behind this screen, and all of it depends on subsea infrastructure that already works. We did not invent putting compute underwater. We changed what the compute is for.

www.subseacloud.com

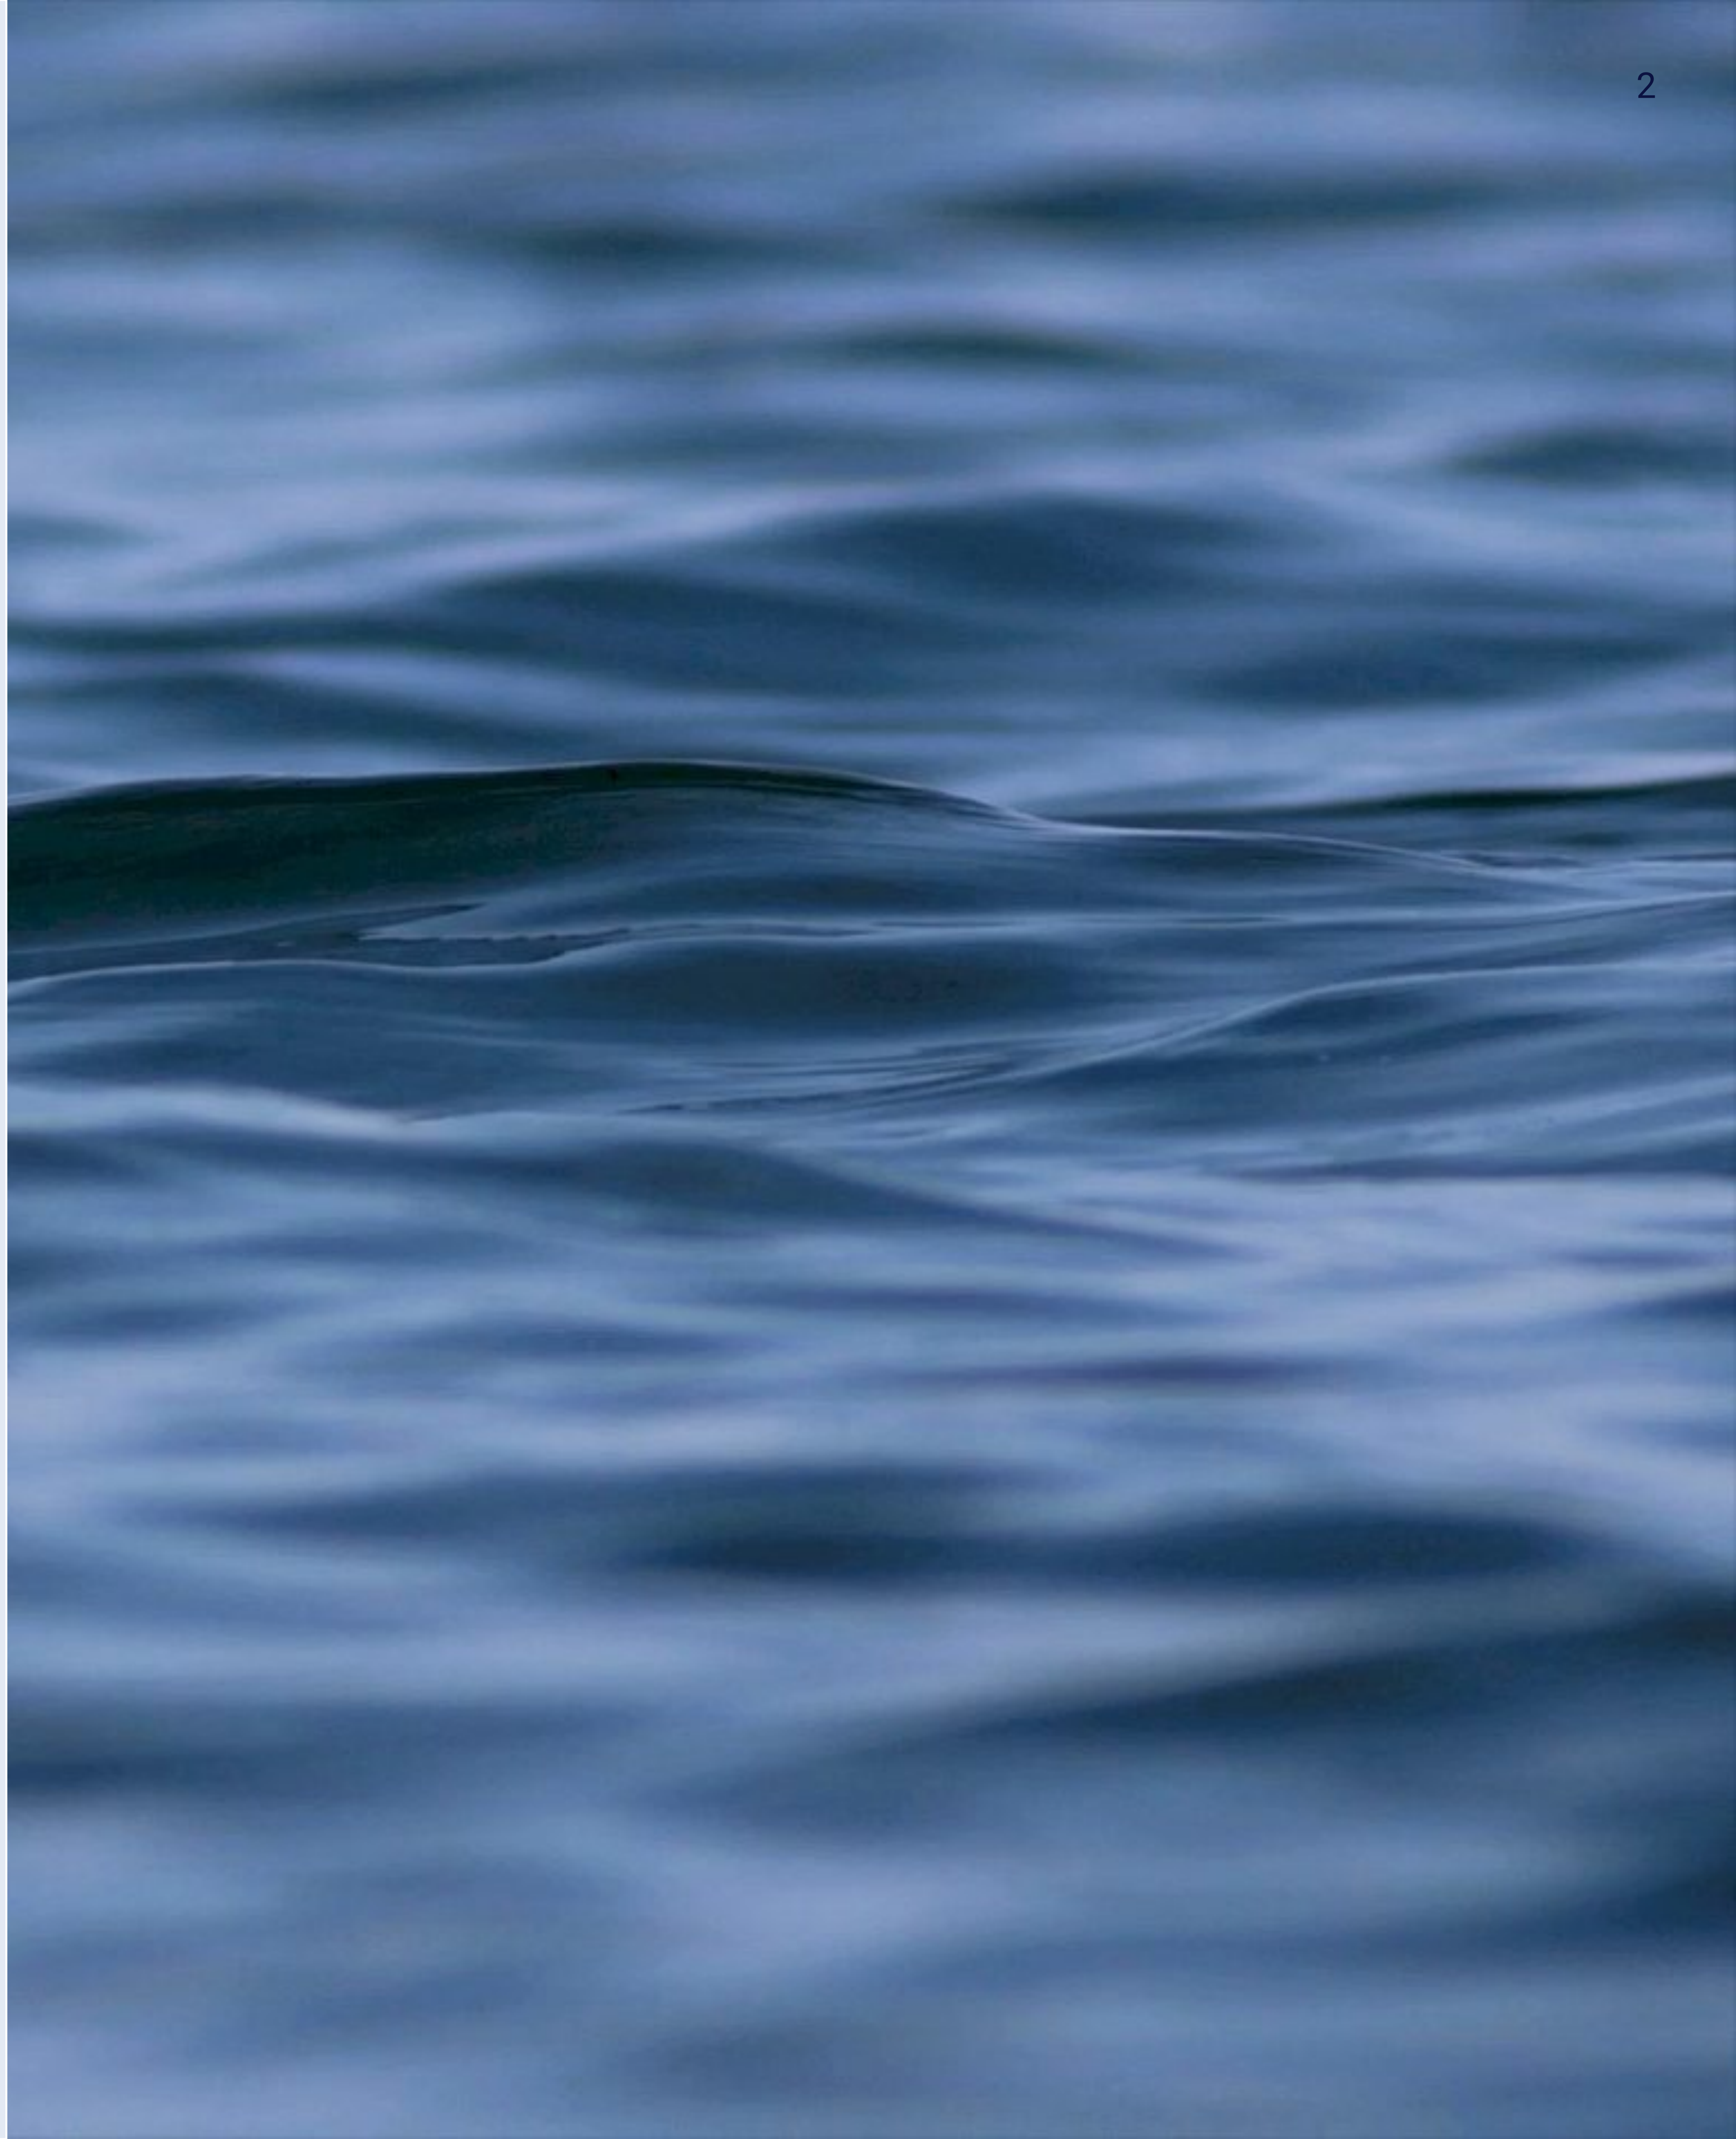

SUBSEA CLOUD

Subsea Data Centers

www.subseacloud.com

2026

THANK YOU

Subsea Data Centers

www.subseacloud.com

2026