Could Underwater Data Centers Pose a Risk to AI Treaty Verification?

Ashmita Rajmohan, James Teague, Yannick Muehlhaeuser

Executive Summary

We expect substantial AI training in underwater data centers (UDCs) will eventually be possible, primarily owing to their potential for extremely superior energy and cooling efficiency. However, UDCs come with many obstacles and downsides, ultimately rendering them an unlikely path to treaty evasion when compared to hypothetical alternatives, such as underground or land-based data centers that are disguised as alternative industrial processes.

For construction, there is a very limited supply of materials and marine capacity for fabricating pressure sealed hulls or for laying subsea cables. This imposes significant economic and operational burdens, which we expect would likely increase when construction must remain covert to avoid detection. Definitively evading optical and SAR satellite surveillance of these ships and fabrication yards would likely require the evading actor to covertly build some additional industrial capacity, operate in an unsuspecting location at times when monitoring may be limited, and then balance this with the incentive to quickly finish construction, the most detectable phase, as soon as possible.

It is not yet operationally feasible to conduct maintenance and hardware updates on a frontier scale UDC, such as via divers or autonomous underwater vehicles. This drastically reduces the overall reliability and efficiency of an underwater cluster over its operational cycle, which sits between 3 and 5 years. However, whilst the probability of detection during the buildout stage is significant, this work suggests that an already operational UDC could feasibly evade detection in the long term, provided it is built in accordance with a relatively specific set of specifications. Thus, while we believe UDCs likely constitute an unlikely path to treaty evasion, we believe there is still non-zero risk that should be addressed by beginning to operationalise the detection modalities discussed.

Introduction

Most [proposals] for international agreements that place restrictions on frontier AI development require robust verification mechanisms to monitor compliance. A critical challenge for such agreements is detecting undeclared or "dark compute facilities" that states might deploy to evade restrictions on frontier model training. Underwater data centers (UDCs) have been proposed as a possible evasion vector. However, the technical feasibility of deploying UDCs at the scale required for training frontier models remains unclear and their detectability has not been seriously investigated.

We examine the state of current UDC deployments (Microsoft Natick, Highlander, Subsea Cloud), evaluate construction complexity relative to traditional data centers, and investigate their feasibility for large-scale AI training. We then assess the detectability of large-scale subsea clusters via thermal, acoustic, and visual surveillance methods.

Overview

Underwater data centers are sealed capsules of computing infrastructure deployed on the seabed, using cold ocean water for cooling in place of conventional cooling systems. This section covers their cooling methods, physical design, operational reliability, and current deployments.

Cooling

Underwater data centers use two primary cooling methods. In the first, servers are immersed in a dielectric fluid that absorbs heat directly from the chips, which is then dissipated to the surrounding seawater through the pod's cooling system (Cao et al., 2022). In the second, cold seawater is piped through heat exchangers positioned behind the server racks (Cao et al., 2022; Cutler and Fowers, 2017). Both approaches eliminate the conventional chillers, cooling towers, and HVAC systems that dominate the energy budgets of land-based facilities, where HVAC has historically accounted for roughly 25 to 40 percent of data center electricity consumption in conventional air-cooled facilities (US DOE FEMP; Shehabi et al., 2016).

At roughly 100 metres depth, seawater temperatures are consistently cold and stable year-round, with peer-reviewed feasibility studies of deep seawater cooling systems estimating annual cooling energy savings of approximately 79 percent relative to conventional chiller-based systems (Kumar et al., 2025).

Physical Design

The basic form factor is a sealed, pressure-resistant steel cylinder designed for multi-year unmanned operation, sometimes referred to as "lights-out" deployment (Microsoft Project Natick). Two main design families have emerged, each pairing a pressure strategy with a corresponding cooling approach. Sealed capsules, as used in Project Natick, maintain one atmosphere of internal pressure and are filled with dry nitrogen, which eliminates the oxygen and humidity that drive corrosion in conventional facilities (Microsoft Project Natick; TechTarget, 2023); heat is moved from the servers to the ocean via fans and seawater piped through internal heat exchangers (Cutler and Fowers, 2017; DCD, 2022). Pressure-equalised pods, as proposed by Subsea Cloud, match internal pressure to the surrounding ocean (The Register, 2022; DCD, 2022). Because gases compress at depth, pressure equalisation requires filling the enclosure with a nearly incompressible dielectric fluid, which doubles as the immersion coolant: servers sit directly in the fluid, and heat transfers through the container walls to the surrounding seawater (The Register, 2022; Brightlio, 2025). The pressure-equalised approach in principle allows for simpler construction and reduced structural requirements, since the walls no longer need to withstand the full pressure differential at depth (DCD, 2022; Brightlio, 2025).

Each capsule is connected to shore by a submarine cable bundle that carries electrical power and fibre-optic data links, essentially functioning as an extension of a land-based network (Microsoft, 2020; GCN, 2020). Deployment and retrieval are handled by crane barges or similar marine vessels. The capsules are lowered to the seabed and anchored in place, and the same process is reversed when a unit reaches end of life or requires servicing (DCD, 2025; Microsoft, 2020).

Reliability

Over its two-year deployment, Natick recorded a server failure rate roughly one eighth that of equivalent land-based facilities, with less than one percent of servers failing compared to approximately eight percent on land (Microsoft, 2020; Tom's Hardware, 2024). This improvement has been attributed to the inert nitrogen atmosphere inside the capsule and the absence of human-induced disruptions such as accidental bumps, cable pulls, and maintenance errors that are common in staffed facilities (Microsoft Project Natick; Microsoft, 2020).

However, because the capsules are sealed and inaccessible during operation, individual server failures could not be repaired in situ like land-based data centers, where full repairs can be completed within hours and training continuity can be restored within minutes. Instead, workloads are redistributed across remaining healthy nodes, and failed hardware stays offline until the capsule is retrieved at the end of its operational cycle (Microsoft, 2020; GCN, 2020). For sustained deployments spanning two years or longer, these individual failures could compound, likely rendering the long-term performance of UDCs progressively worse than equivalently provisioned land-based clusters; at least until a reliable mechanism for conducting repairs whilst pods remain operational is found.

Current and Past Examples

Microsoft's Project Natick was the first underwater data center programme and ran in two phases. Phase 1 deployed a single rack off the California coast from August to December 2015. Phase 2, a full-scale 864-server module, operated on the seabed off Orkney from June 2018 to July 2020. Owing to difficulties in servicing, upgrading, and replacing hardware, which would be vital for taking the concept into full production, Microsoft formally ended the project in June 2024.

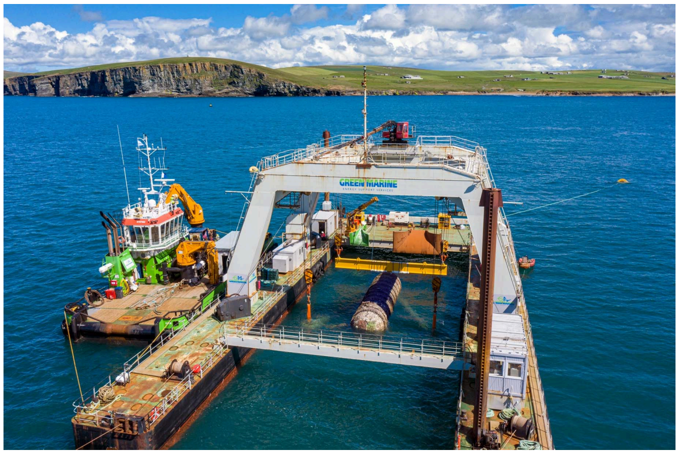

Green Marine's barge raises Microsoft's Project Natick pod off Orkney, Scotland, in 2020, after its two-year underwater trial. The cylinder, coated in marine growth, hangs beneath the gantry crane.

A worker pressure-washes barnacles and algae from the recovered Natick pod, revealing the Microsoft logo beneath two years of biofouling.

Highlander (via subsidiary HiCloud) is the only company currently operating commercial underwater data centers. Its Hainan deployment has two capsules in the water, with Module 2 going live in February 2025, and its Shanghai facility completed Phase 1 construction in October 2025, making it the first UDC directly linked to an offshore wind farm.

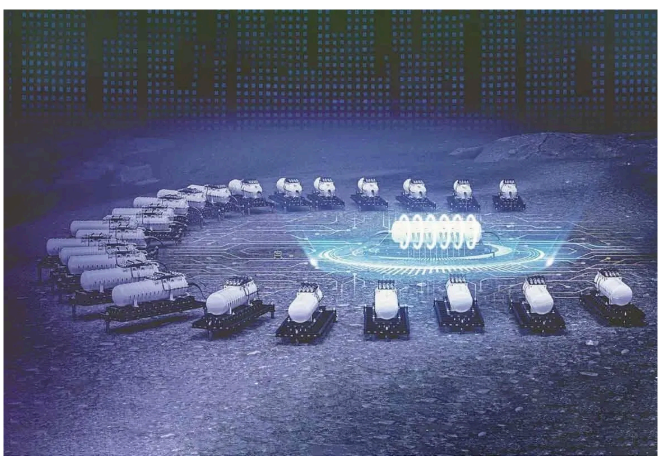

A HiCloud rendering of its underwater data center concept, with cylindrical server pods arrayed across the seabed around a central hub.

A HiCloud visualization of submerged data center pods arranged in an arc on the ocean floor, illustrating its modular subsea deployment off China's coast.

Subsea Cloud announced Project OTTO, a rolling trial off South-West Norway, with its first phase scheduled for October 2024, though independent verification of deployment outcomes is limited. Its Jules Verne pod near Port Angeles, Washington was announced as a demonstration unit.

NetworkOcean, a YC-backed startup, has announced a 0.5 MW test capsule for San Francisco Bay but has faced regulatory obstacles, with at least two California agencies flagging the lack of required permits. Their plans are to submerge the capsule a few metres below the surface for roughly an hour, with the company claiming the deployment would eliminate on-site water consumption and reduce power usage by up to 30% through passive seawater cooling.

Spec	MS Natick Ph1 (2015)	MS Natick Ph2 (2018–20)	Highlander Hainan (2022 pilot, 2025)	Highlander Shanghai (wind powered, 2025)	Subsea Cloud Jules Verne	NetworkOcean (pre-deploy)
Dimensions	~3m × 2.1m MS Natick	12.2m × 2.8m; ~40ft pod MS Natick	~18m × 3.6m ECNS	Similar modular design; not detailed Xinhua	~20ft x 8.5ft (6m) container Subsea Cloud	Unknown
Depth	~11m	35.7m	~35m Int. Eng.	Depth unspecified	~9m DCD	"Just under surface" DCD
Servers / Racks per pod	1 rack, ~24 servers, 1 pod	12 racks, 864 servers, 1 pod	24 racks, 400 servers per pod; 2 pods Dev. Telecoms	198 racks (Phase 1), 2 pods Sci. Am.	16 racks, ~800 servers/pod	1 test capsule planned YC
CPU / Accelerator used	Standard Azure hardware (likely Xeon)	Xeon E5 v3/v4 + Arria 10 FPGAs; no GPUs.	Undisclosed	Undisclosed	Undisclosed	NVIDIA H100 (2,048 reserved); advertises NVL72/GB200.
Total power	Low (test-scale)	240 kW	Unclear	Ph1: 2.3 MW (2 pods); barge target: 24 MW Tom's HW	0.5–1.5 MW/pod Subsea Cloud	0.5–1 MW test capsule; barges scalable to 200 MW+.
Compute (FLOPS)	~tens TFLOPS	~100s TFLOPS	No benchmark. Planned: 2 exaFLOPS in a 100-module buildout Yicai	Undisclosed	Undisclosed	Planned barge: ~2,026 PFLOPS FP16 in a full buildout (2,048 × H100). NVIDIA H100 spec (989 TFLOPS/card, not confirmed underwater)
Power source	Shore grid	Shore grid	Shore grid	Direct offshore wind via submarine cable (>95%); onshore grid backup (<5%). Int. Eng.	Not specified.	Not specified.

Construction and Maintenance Complexity

Building an underwater data center draws on a different industrial base than building one on land. A conventional facility requires concrete, structural steel, and HVAC contractors — industries with tens of thousands of qualified firms worldwide. An underwater facility requires pressure vessel fabrication, submarine cable installation, and marine engineering capability and faces severe limitations in maintenance, hardware refresh and potential vulnerability to sabotage.

Hull fabrication. Currently, the core structural element of an underwater data center is a sealed, corrosion-resistant steel pressure vessel. According to the naval analyst H I Sutton, approximately 17–18 countries currently possess the capability to design and build full-sized submarines.[1] This asymmetry makes pressure vessel fabrication a strong supply chain chokepoint: the activity is concentrated in a small number of known, large shipyards and is inherently visible to satellite observation. However, this bottleneck may not be permanent. With significant investments in scaling pressure vessel fabrications, the capsules may become factory-built modular units built in serial production in multiple yards, eroding both the cost premium and the monitoring advantage of supply chain concentration.

Submarine cabling. Every underwater data center requires at least one submarine cable bundle carrying both electrical power and fibre-optic data links from shore. Installing these cables requires specialised cable-laying vessels. According to DCD, the global fleet of ships capable of laying and maintaining submarine cables numbers only about 60.[2] The major operators are a known and small set of companies[3], and their vessels are tracked via AIS.

Marine operations. Activities such as cable laying, seabed preparation, pod deployment via crane barge, and umbilical connection typically account for 40 to 60 percent of total underwater data center project cost. This is the single largest cost premium over land-based construction, where site preparation and utility connection, while slow, draw on abundant and inexpensive civil engineering resources.

Maintenance. In the current generation of UDCs, hardware failures such as a dead GPU or a failed network switch can't be repaired until the entire pod is retrieved. Subsea Cloud claims retrieval in 4–16 hours for its lighter, pressure-equalised pods. While this has not been independently demonstrated at scale, it appears that these are realistic numbers after significant time has been invested in development.

---

[1] Major yards are in China (CSSC/CSIC, including Wuchang), France (Naval Group), South Korea (Hanwha Ocean, HD Hyundai), Japan (Mitsubishi, Kawasaki), Germany (ThyssenKrupp Marine Systems), Sweden (Saab Kockums), and several others including India, Italy, and Turkey

[2] Not all 60 are equivalent: many are smaller repair vessels or fibre-only ships. Power cable installation — which is what an underwater data center needs — requires larger vessels with higher-capacity carousels and tensioners, further narrowing the available fleet to perhaps 20–30 ships.

[3] Prysmian, SubCom, Nexans, NEC, Global Marine, Jan De Nul, SBSS

Hardware refresh. The hardware refresh cycle is widely cited as the most significant barrier to underwater data center economics. GPU generations evolve every 12–18 months, and sealed pods designed for 3–5 year deployment cycles cannot accommodate mid-cycle upgrades. This was a central reason Microsoft discontinued Project Natick: sealed pods are incompatible with the rapid hardware iteration that hyperscale cloud operators require for commercial competitiveness.

Sabotage. A final point is the vulnerability to physical sabotage. A 2024 study showed that off-the-shelf underwater speakers tuned to the hard-drive resonant band (~300–1,300 Hz) could crash server networks by inducing resonant vibration in hard drives, degrading or crashing storage nodes within minutes. This presents an additional vulnerability alongside other kinetic forms of sabotage.

Frontier AI training feasibility

Power delivery and cooling for such a training run appears quite feasible. 100k H100s draw ~70 MW at the GPUs alone, but server overhead (CPUs, NICs, PSUs) and switch fabric push this toward ~150 MW. A frontier cluster in the relevant 3–5 year window may also require denser hardware (H200s or B200s), suggesting 200 MW as a plausible high-end. Submarine HVDC power cables routinely carry GW-scale loads, so delivering 150-200 MW via a dedicated submarine power cable is well within proven engineering. Cooling is where UDCs generally have an advantage over land-based clusters, so this does not constitute a bottleneck.

Interconnect is the most plausible bottleneck. For one, the network fabric itself is physically massive. Beyond the GPU racks, a 100k H100 equivalent cluster also needs hundreds of leaf switches and dozens of spine switches, each consuming tens of kilowatts, each requiring precisely structured cabling to every other switch. In a reference 8,192-GPU design, 256 leaf switches connect to 36 spine switches housed in dedicated spine racks. This means substantial additional engineering effort beyond what UDCs are currently designed for and likely increases the required power delivery (accounted for above). While it may not be economical, this appears solvable for a motivated actor.

Cable distance is another aspect of networking that could constitute a bottleneck. Between racks, direct-attach copper works up to about 3m, active optical cables up to 30m, and multimode fibre up to roughly 50m for the leaf layer. This means the physical layout must be extremely dense and precisely arranged. This likely means additional engineering effort and risk.

The most significant obstacle from the networking side appears to be the required maintenance. Common maintenance tasks such as unplugging and plugging back in the InfiniBand transceiver are fundamentally hands-on operations.[4] Meta's Llama 3 405B training run provides

[4] SemiAnalysis describes the common operational issues in GPU clusters as "flapping IB transceivers, GPUs falling off the bus, GPU HBM errors, and silent data corruption," and that resolution often requires "unplugging and plugging back in the InfiniBand transceiver or cleaning the dust off of fiber cables."

the most detailed public data point: over 54 days on 16,384 H100 GPUs, the cluster experienced 419 unexpected interruptions — one every three hours. Optical failures, which are arguably the most severe reliability problem, also don't get reduced by a less corrosive atmosphere the way chip failures themselves are. This means that conducting a large training run in an UDC will either require very difficult underwater maintenance or massive redundancy to route connections around points of failure.

Overall, interconnect appears to constitute the most significant obstacle to training a large AI model in an UDC. The limitations outlined could potentially be overcome by a motivated state actor that is willing to spend large amounts of money. Thus, if dark compute via UDC is a concern, it depends on the probability of detection and the feasibility of alternative pathways.

Agreement Evasion Feasibility

The detection modalities most applicable to operational UDCs are thermal infrared (TIR) imaging and acoustic surveillance, whilst optical and synthetic aperture radar (SAR) imaging are capable of detecting the logistics and infrastructure activity associated with building and installing a UDC cluster. The effectiveness of TIR imaging is likely negligible given the effects of heat dispersion in coastal waters; acoustic detection is physically plausible but faces significant operational bottlenecks; and optical and SAR monitoring are most powerful during the deployment window.

Thermal Detection

A frontier-scale UDC cluster drawing ~150–200 MW would continuously reject waste heat into the surrounding water column through seawater heat exchangers or immersion fluid systems. The relevant question is whether this produces a detectable sea surface temperature (SST) anomaly. At the ~35m deployment depths of current commercial UDCs, and given that frontier-scale training would likely require an area on the order of ~1.7–9 hectares[5], detection is unlikely in the majority of open coastal conditions, where current velocity usually sits well above 0.17m/s. Waste heat expelled at this depth would need to rise through a coastal water column before reaching the surface, implying an SST change below standard deviation (± 0.3 – 0.5°C in the lower variability of tropical ocean basins).[6]

Conventional Thermal Infra-Red (TIR) imaging satellites, such as Landsat-8, offer coarse resolution (100m) relative to the scale of UDCs likely thermal signature. Whilst Ahn et al. (2006) demonstrated that satellite thermal-infrared imagery (NOAA AVHRR and Landsat) could track the warm-water plume from Korea's ~5,900 MW Younggwang nuclear plant up to 100 km from the outfall, this discharge enters the sea from a concentrated point at the surface rather than being dispersed at depth. A UDC's smaller, deeper discharge would not produce a comparable

[5] See Appendix A.5
[6] See Appendix A.6

surface signature. This limitation persists even for newer commercial TIR instruments. Even though SatVu's HotSat-2 can provide day-and-night coverage at a 3.5m resolution, it would still be unable to detect an SST anomaly that dispersion has already diluted below natural variability. Independent of resolution, infrared sensors are also limited by cloud cover, creating persistent monitoring gaps.

The three key variables affecting detectability are therefore depth, dispersion and location. Deeper and more widely dispersed pods will reduce the thermal signal at the sea surface, as discharge will rise through a more stratified water column and over across a greater volume. The additional costs and engineering constraints of deeper and more dispersed pods can be offset by building UDCs in waters with an increased current velocity: significantly reducing the likelihood of a detectable SST anomaly from shallower or more dense clusters. This effectively rules out UDC construction in calmer waters[7] (e.g. coves, sheltered inlets, or inland waters like lakes) where covert construction and operations may initially appear simpler. Operators seeking to evade thermal detection would therefore have strong incentives to deploy as deep as structurally and economically feasible, as distributed as interconnect constraints allow, and in as turbulent waters as a hull can feasibly stand without affecting efficiency.

Acoustic Detection

If both basic UDC designs were deployed at the scale required for frontier training, they would each generate an acoustic signature. For the pressurised capsule design demonstrated by Project Natick, the fans and internal heat exchangers would produce broadband noise with blade-rate tones (roughly 20–150 Hz), and the power-conversion transformers add a stable tone at 100 or 120 Hz (twice the mains supply frequency), both transmitted through the hull into the surrounding water. Dielectric-immersion pods, such as Subsea Cloud's, lack that fan and pump noise, but their internal sources sit in a fluid whose acoustic impedance is much closer to seawater than the gas-to-steel boundary of a sealed hull, so the transformer tone would likely couple into the water more efficiently. Despite the two designs likely differing in their dominant source; an order of magnitude scaling estimate plausibly places their total acoustic signature at roughly 115–130 dB re 1 µPa at 1 m, with substantial uncertainty.[8]

Unlike submarines, which are moving targets that use active noise-reduction engineering, UDCs are stationary, continuously operating industrial facilities, suggesting they may be easier to track via traditional acoustic surveillance systems. However, the US Navy's primary detection mechanism for submarines is the SOSUS network: arrays of sea-bed mounted hydrophones at the axis of the SOFAR channel, which whilst capable of passively detecting acoustic power of less than one watt at ranges of several hundred kilometres, operates at 600–1,200m depths. Current UDC deployments at 30–35m sit in shallow coastal waters, where propagation losses are substantially higher and bottom reverberation masks nearby targets.

---

[7] See Appendix A.7
[8] See Appendix A.8

Without the monitored state's cooperation on naval monitoring, effective detection of a frontier-scale UDC would likely require the covert deployment of acoustic sensors. Existing civilian infrastructure, such as oceanographic moorings or research hydrophone arrays, could provide data without active sensor deployments, but coverage near specific coastal sites cannot be assumed and would require deliberate pre-positioning under scientific cover. A purpose-built sensor such as the AN/SSQ-53 DIFAR sonobuoy can operate for up to 8 hours by transmitting data back by radio link, covers the transformer frequency range and harmonics, and can provide directional bearings to a source: such that two sonobuoys can triangulate a contact. Given the source level estimated above, and the transmission losses characteristic of 35m underwater depths, a sonobuoy would need to be positioned within 0.5-5km of the cluster to achieve detection.[9]

The practical implication is that passive acoustic monitoring requires sensors inside the target state's 12 nautical mile (22km) territorial sea. They could be deployed from a covert submarine, lowered from a surface vessel under civilian cover, or dropped from altitudes up to 30,000 feet. However, each method carries significant risk: a state concealing a UDC would plausibly monitor the surrounding water and likely approach routes for intruding submarines and suspicious foreign objects, and track anomalous surface vessel activity using satellite AIS and coastal radar.

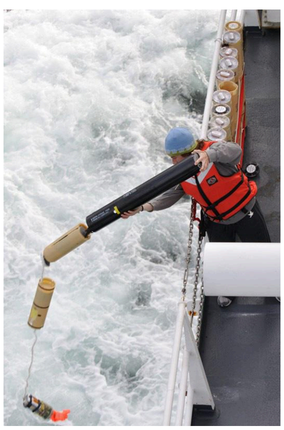

A sonobuoy is launched from the NOAA Ship Oscar Dyson. Image courtesy of NOAA.

A final consideration for covert detection is the use of active sonar. Small Autonomous Underwater Vehicles (AUVs, such as REMUS 100) equipped with side-scan sonar can image a swath of seabed roughly 150m wide in a single pass, so a cluster spanning ~1.7–9 hectares could be surveyed quickly. The use of AUVs has some advantages over sonobuoys, such as outputs being a clear image rather than an inferred contact, and much more efficiency in space covered over time. Against an alert host however, this is not reliably stealthy, since active sonar pings the seabed, the vehicle's noise is highly detectable by purpose-built intruder sonar, and its launch and recovery expose deployers to the same surveillance as any other vessel.

Acoustic detection is therefore possible in principle, but its feasibility is heavily dependent on the specific geography and the counterintelligence posture of the monitored state. With that said, the more aggressive the countermeasures a state deploys (e.g. restricting fishing vessels, imposing exclusion zones, or deploying acoustic decoys), or the more restrictive a state is over the permitted monitoring regimes the treaty can permit, the more suspicion those measures themselves attract. For example, if a state rejected a treaty clause mandating the bilateral deployment of seabed

[9] See Appendix A.9

hydrophone networks or regular AUV surveys (akin to IAEA inspection protocols) this would generate significant suspicions of concealment effort.

Optical and SAR Imaging Satellites

Neither optical satellites nor synthetic aperture radar (SAR) can detect an operational UDC directly. Optical sensors require reflected sunlight, which attenuates rapidly in seawater from 30-35m depths. SAR uses targeted microwave pulses which, while capable of penetrating cloud cover and operating at night, produce near-zero radar return from submerged objects. While SAR techniques have been used to detect bottom topography and depth gradients via surface wave modulation, an UDC cluster would likely not produce an SAR detectable perturbation in the local wave field.

Optical and SAR imaging are nonetheless relevant as deployment-phase detection tools. As noted earlier, there are approximately 60 ships worldwide capable of laying submarine cables, with only a small subset being capable of deploying the power-cables required to deliver hundreds of megawatts to a seabed cluster. SAR can be used to identify and classify specific vessel types, which can be fused with satellite Automatic Identification System (AIS) data to flag vessels that have disabled or modified their transponders; which is a technique already used in sanctions enforcement and fishing regulations. Similarly, optical and SAR satellites have been used in open-source intelligence contexts to monitor shipyard activity, track construction progress, and infer production rates. Assuming the pool of shipyards capable of manufacturing pressure vessels for UDCs is 20–30 globally, the same methodology would allow monitors to track abnormal production volumes or unscheduled deployments.

The key limitation for optical and SAR monitoring is timing: a cable installation completed before a verification regime came into force, or conducted in a manner consistent with other commercial activities, would not generate a distinguishing signature. However, given the maintenance required for frontier training that was outlined earlier, it is likely that a recurring surface signature would be generated and detected throughout a cluster's operational cycle.

Synthesis

Taken together, these three modalities suggest that an operational, frontier-scale UDC could potentially be concealed completely; however, it is likely to be detected during construction or maintenance if SAR and optical satellites are actively used for surveillance of the cable-laying vessel fleet and fabrication yards. Whilst thermal detection is unlikely to be viable for anything other than operational UDCs in very shallow or calm waters (e.g. lakes or coves), acoustic detection could function as a verification mechanism, owing both to the range of counterintelligence measures actors must deploy to combat small sensor deployments, and to the restrictions a state would have to place on permitted monitoring under the treaty.

Conclusion

Underwater data centers are technically feasible as a dark compute pathway for 100,000 H100 equivalent training runs, but only under demanding and strategically distinctive conditions. The predominant obstacle seems to be the large amount of hands-on maintenance that is currently required during larger AI training runs which is incredibly difficult under water. A state actor willing to absorb significant cost and timeline premiums could plausibly overcome this, but the resulting facility is strategically uneconomic — one that makes sense only if the operator's primary objective is concealment rather than cost or performance efficiency. One important implication of this is that UDCs built for commercial purposes and UDCs built for evasion have different engineering signatures.

No single detection modality forecloses the UDC evasion pathway, but a layered approach raises the bar substantially. Deployment-phase surveillance of the small pressure-vessel fabrication pool and the AIS-tracked power-cable-laying vessel fleet are the most effective detection pathways. If working in conjunction with supply chain focussed measures, satellite imagery is likely powerful at tracking UDC buildout. Detecting a UDC once it has been built will be much harder, but the maintenance might still be clearly detectable. Hydrophones and AUVs provide a promising pathway to detection, especially if parties to an agreement mutually agree to host them in areas where UDCs may possibly be built. Overall, the detection methodologies discussed work best in a layered approach, where one layer compensates for the weaknesses of another.

How feasible a path to dark compute UDC presents depends in large part on the available alternatives. Underground data centers — concealed in mines or excavated cavities, with off-grid power and construction disguised as civil or mining activity — offer substantially greater ease of hardware maintenance. Concealment of data centers as ordinary industrial buildings presents an even lower barrier: equipment can be delivered in ordinary shipments, personnel movements are unremarkable, and the construction signature is minimal compared to the marine logistics operation that any frontier-scale UDC requires. Compared with UDCs, these data centers may be easier to detect after they have been built, but this is likely compensated for by the greater difficulty of hiding UDC construction.

Given the obstacles outlined and the feasibility of detection during the buildout stage, we do believe, while in principle possible, UDCs likely constitute a much harder path to treaty evasion compared with the alternatives. Nevertheless we recommend that the intelligence agencies of nation states actively monitor developments in the area and research detection methodologies.

Acknowledgements

This work was done as part of the Orion AI Governance Initiative by Arcadia Impact. We thank Joe Hardie for general support and comments. Further input was provided by Christina Krawec, Joshua Turner, Mauricio Baker, Ben Harack, Mehmet Sencan, James Petrie, Robi Rahman and Aaron Scher.

Appendix A: Supporting Calculations

This appendix contains the full derivations for the order-of-magnitude estimates cited in the main text, retained here for readers who want to verify or scrutinize the underlying assumptions.

A.5

Order-of-magnitude estimate. Assuming 150-200 MW required power, and 2 MW per pod (feasible given demonstrated designs) ≈ 75–100 pods. Assume 15–30 m centre-to-centre spacing between pods (for deployment logistics and/or heat-discharge flow), i.e. 225–900 m^2 total space needed per pod. 75 × 225m^2 ≈ 1.7ha; 100 × 900m^2 ≈ 9ha.

A.6

In the most extreme case: 200 MW heat rejected ÷ (4.0×10^6 J m^{-3} K^{-1} [seawater heat capacity] × 0.3°C [maximum covert SST anomaly] × 100m [even lower than the densest feasible cluster width] × 10m mixing depth [under summer stratification]) ≈ 0.17 m/s as the lowest current velocity viable for evasion.

A.7

In the most conservative case: 150 MW ÷ (4.18×10^6 J $m^{-3}K^{-1}$ [freshwater heat capacity] × 0.005m/s [wind-driven circulation] × 300m [widest feasible cluster width] × 35m [full water column mixing]) ≈ 0.68°C (rising to ~2.4°C under summer stratification, where mixing depth is limited to 10m).

A.8

This is an order-of-magnitude estimate (±15dB). The average audible sound levels for liquid-immersed network transformers and step-voltage regulators in the 1,001–2,000 kVA range are 60-61dB. Converting an airborne SPL of ~60–61 dB re 20 μPa to an equal-intensity underwater pressure equivalent gives ~121–123 dB re 1 μPa, after adding 26 dB for the pressure-reference change and ~35–36 dB for the air/water acoustic-impedance difference. Subtracting ~15-25dB for structural hull coupling, and adding ~20dB for incoherent summation across the 75-100 pods yields a source level of ~115–130dB re 1μPa at 1m. For reference, the source level of an offshore wind turbine can be approximated as 156dB re 1μPa at 1m, with tones below 1 kHz. Without the same MW-scale rotating machinery, UDCs would plausibly sit under 130dB.

A.9

Detection requires the received level (source minus transmission loss) to exceed the ambient-noise level in the transformers approximate frequency. At ~100 Hz a shipping-dominated coastal floor is 75-80 dB re 1 μPa^2/Hz (Hildebrand, 2009, p11); integrating the transformer tone over a narrow ~1 Hz band and applying modest narrowband processing gain gives an effective detection threshold of ~75 dB re 1 μPa. With a 115 –130 dB source, shallow-water three-halves spreading at 100 Hz (TL ≈ 15 $\log_{10}$ r) and negligible absorption (~0.02 dB/km), the maximum tolerable transmission loss is ~40–55 dB, giving a detection range of ~0.5–5km.