Surveying the Evolution from Offshore Data Centers to Autonomous Oceanic Computing

Benedetta Picano[a,*], Dusit Niyato[b]

[a]University of Florence, Via di Santa Marta, Florence, 50139, Italy,
[b]Nanyang Technological University, Nanyang Avenue, 639798 Singapore,

ARTICLE INFO

Keywords:
Offshore data centers
Oceanic computing
Autonomous marine platforms
Sustainable AI infrastructure

ABSTRACT

The rapid growth of artificial intelligence and cloud services is increasing demand for data centers. However, at the same time, suitable land, grid capacity and cooling water are becoming constrained. Moving computing infrastructure offshore offers a different approach. First, marine environments provide access to renewable energy. Moreover, they provide abundant space and seawater as a natural heat sink. This idea is already moving beyond experimentation, with submerged data-center pods demonstrated at scale, floating and vessel-based facilities under development, and autonomous wave-powered computing platforms attracting commercial investment. Nonetheless, these developments have largely been studied separately. In this paper, we introduce *oceanic computing* as a unifying concept for computing infrastructure that uses marine resources and operates with less or without human intervention. We organize this emerging design space along four infrastructure dimensions and examine six operational domains: energy harvesting, workload management, thermal control, communication, maintenance, and positioning. We further show several key design principles. They include the use of computational demand as a control variable for energy conversion. These principles appeared in patents well before they entered the academic literature. Bringing these developments together reveals that the six operational domains cannot be managed independently. Instead, they form a coupled resource-allocation and control problem. We identify positioning as a particularly unexplored dimension and discuss the legal, environmental, and economic questions that must be resolved before oceanic computing can develop into a viable computing infrastructure.

1. Introduction

The expansion of artificial intelligence, cloud computing and data-intensive digital services is driving a sustained increase in computing, storage and networking demand [1]. This growth is placing increasing pressure on the electricity, water and land resources needed to support data centers. Global estimates of data-center electricity use continue to rise as workloads, hardware and deployment models evolve [2]. Data centers consumed about 415 TWh of electricity in 2024, equivalent to around 1.5% of global electricity use, and their consumption has grown by roughly 12% per year since 2017—more than four times the growth rate of overall electricity demand. By 2030, consumption is projected to more than double to approximately 945 TWh, or almost 3% of the global total [3, 4]. Although around half of this additional demand is expected to be supplied by renewable generation [3], data-center growth is highly concentrated geographically. The United States accounted for 45% of global data-center electricity consumption in 2024, China for 25% and Europe for 15%, with nearly half of US capacity located in only five regional clusters [3]. Water availability is becoming an additional constraint in many regions [5]. Therefore, these pressures are making the location and design of future computing infrastructure an increasingly important question.

Improving the sustainability of data centres therefore requires more than making servers themselves more efficient. Computing resources, power supply, cooling infrastructure, and workloads must be designed and managed as a coupled system [6, 7]. Among these components, cooling is particularly important as it represents one of the largest non-IT energy loads in modern facilities [8, 9]. Conventional air cooling becomes increasingly inefficient as rack power densities rise, prompting a shift towards approaches such as direct liquid cooling and free cooling. Direct liquid cooling removes heat close to the processors, whereas free cooling uses ambient air or water when external conditions allow mechanical chillers to be reduced or switched off [8, 9]. Depending on the climate and cooling configuration, such approaches can produce substantial energy savings, with reductions of up to 67.2% reported relative to conventional cooling [8]. Their effectiveness, however, remains tied to local environmental conditions. More fundamentally, even highly efficient cooling does not remove the dependence of a data center on land, grid connections, and local water resources. These constraints become especially important in densely populated coastal regions and city states, where available land is scarce, but marine space and renewable-energy resources remain comparatively underused [10].

Marine environments offer several resources that are difficult to obtain simultaneously on land. Offshore wind, wave, tidal, solar and hybrid systems can provide renewable electricity [10, 11], while seawater offers a large natural heat sink that can reduce the energy required for cooling [1, 11]. Moving infrastructure offshore can also ease pressure on scarce land and enable modular expansion near densely populated coastal regions [1]. These advantages become

*Corresponding author
benedetta.picano@unifi.it (B. Picano)
ORCID(s): 0000-0003-4970-1361 (B. Picano)

particularly important where conventional data-center development is constrained by geography. A system-level study comparing onshore, nearshore, and offshore data centers, for example, considered location, renewable-energy supply, mobility, cooling, cost, environmental impact and scalability as interconnected design choices. It showed that offshore configurations can become attractive in land-constrained environments such as Singapore, particularly when data-center expansion is coordinated with offshore renewable-energy development [1, 12]. This perspective suggests that the ocean should not be viewed simply as a new location for data centers, but as an environment in which computing, energy, and cooling infrastructure can be designed together.

The idea of placing computing infrastructure at sea emerged much earlier. A patent filed in 2007 and granted in 2009 already described a floating data center composed of crane-removable computing modules, an array of wave-energy generators, and seawater-to-freshwater heat exchangers for cooling [13]. Modularity, marine renewable energy, seawater cooling, and power aggregation were therefore integrated in a single architecture nearly two decades ago. What was largely absent, however, was autonomy. The computing equipment consumes the electricity produced by the wave-energy system, but computational demand does not influence how that energy is generated or managed. Questions such as how the platform should adapt its workload to changing sea conditions, how it should communicate with shore, and how it should operate with limited or without human intervention remained open. In this sense, the early concept was essentially only a conventional data center relocated onto the ocean.

The submerged branch of this idea developed first. Microsoft's Project Natick [14], for example, deployed a sealed module containing 864 servers at a depth of about 36 m off the Orkney Islands between 2018 and 2020. The project reported a server failure rate approximately one-eighth of that observed in comparable land-based facilities, attributed primarily to the controlled nitrogen atmosphere and the absence of human intervention [14]. This result demonstrated that removing people from the operating environment can improve reliability. Nonetheless, it also highlighted a broader question: *if computing infrastructure is placed increasingly far from routine human access, how much of its operation must eventually become autonomous?*

In this survey, we define an autonomous offshore data center as marine computing infrastructure that can sense its physical and operational state, anticipate relevant changes [15, 16]. Such infrastructure can coordinate decisions across computing, energy, cooling, communication, maintenance and positioning [17, 18], with limited or without continuous human intervention [14]. Autonomy therefore means more than simply operating without personnel on site. It requires sensing, prediction, decision-making and cyber-physical control to work together so that the facility can respond to changing workloads and environmental conditions. We use the term oceanic computing for the broader paradigm that emerges when such autonomous facilities can be deployed according to marine energy availability, either individually or as coordinated fleets. Autonomy makes an individual offshore platform viable. Accordingly, oceanic computing extends this idea to a distributed computing infrastructure whose capacity can be placed, and in some configurations, moved, according to changing marine resources. This meaning is distinct from initiatives that apply computing and data analytics to ocean science [19]. Here, the ocean is not the object being computed; it is the environment in which computation takes place.

A core feature of oceanic computing is that it brings together two forms of intelligence. First, digital intelligence manages workloads, servers, storage, networking and service requirements. Second, physical intelligence manages the marine platform, its energy resources, structural condition, interaction with the environment and physical position. These two forms of intelligence must be designed jointly. A workload scheduler that ignores renewable-energy availability may request more computing power than that the platform can sustainably provide. Likewise, an energy-management system that ignores service deadlines may preserve stored energy while delaying critical workloads. Repositioning a platform may improve access to wave or wind energy but simultaneously reduce communication quality or consume additional propulsion energy. Similarly, a predicted hardware failure may trigger workload migration, thermal reconfiguration, remote inspection, or even movement towards a maintenance facility. Oceanic computing therefore turns interactions that are secondary in terrestrial data centers into fundamental operating decisions.

Mobility is an important factor of this broader autonomy. However, it is not a requirement for every oceanic computing system. Some facilities may remain fixed on the seabed or moored at the surface while autonomously managing their computing, energy, cooling and maintenance functions [14]. Others may support station keeping, orientation control or limited movement within an operating region [17]. Relocatable platforms could move less frequently in response to severe weather, maintenance needs, changes in connectivity or long-term variations in renewable-energy availability. Recent work has also shown that marine heat waves can reduce the cooling margin available to submerged facilities and has proposed relocation as a means of maintaining operation [20]. At the other extreme, fully mobile platforms would make geographical position itself a controllable resource, jointly optimized with computing demand and environmental conditions. However, such flexibility comes with costs including propulsion energy, structural loading, operational risk, communication disruption and additional capital and maintenance requirements.

These interactions mean that the transition from offshore data centers to autonomous oceanic computing creates a tightly coupled control problem. The platform must decide which workloads to admit, execute, postpone, reduce, checkpoint or migrate. At the same time, it must divide limited energy among computing, cooling, communication, storage,

Table 1
Coverage of the reviewed literature with respect to the dimensions along which this survey contributes. ✓: covered; (✓): partially covered; ✗: not covered.

Work	Facility classes: Submerged pod	Floating / moored	Self-propelled	Analysis dimensions: Off-grid operation	Energy–compute coupling	Positioning	Patent record	Cross-frontier
[6]	✗	✗	✗	✗	(✓)	✗	✗	✗
[7]	✗	✗	✗	✗	(✓)	✗	✗	✗
[8]	✗	✗	✗	✗	✗	✗	✗	✗
[9]	✗	✗	✗	✗	✗	✗	✗	✗
[23]	✓	(✓)	✗	✗	(✓)	✗	✗	✗
[25]	✓	✗	✗	(✓)	(✓)	✗	✗	✗
[15]	✗	(✓)	(✓)	✗	✗	✓	✗	✗
[16]	✗	(✓)	(✓)	✗	✗	(✓)	✗	✗
[21]	✗	✗	✗	✗	✗	✗	✗	✗
[26]	✓	✓	✗	✗	✗	(✓)	✗	✗
[27]	✗	✗	✗	✓	(✓)	✗	✗	✓
[28]	✗	✗	✗	✗	(✓)	✗	✗	✗
[4]	✗	✗	✗	✗	✓	✗	✗	✗
This survey	✓	✓	✓	✓	✓	✓	✓	✓

auxiliary systems and, where applicable, propulsion. Communication must adapt to changing link availability, latency, bandwidth, energy cost and application requirements [21]. The system must also detect degradation, anticipate failures and plan maintenance despite limited physical access. These tasks become particularly challenging in the marine environment, where biofouling and contaminant accumulation progressively change the thermal and structural properties of submerged surfaces [22]. If mobility is available, the platform faces an additional decision: whether the benefits of moving for energy harvesting, cooling, connectivity, safety or maintenance justify the energy and operational costs of doing so.

Research relevant to these problems already exists, but it remains fragmented across several communities. Data-centre research provides mature methods for energy-efficient computing, workload scheduling, thermal management and renewable-energy integration [6, 7]. Underwater data-centre studies examine submerged deployment, cooling, reliability and remote operation [23]. Maritime engineering contributes methods for structural monitoring, navigation, station keeping, digital twins and predictive maintenance [15, 16]. Meanwhile, offshore-energy research develops technologies for marine power generation, storage and operation under uncertain environmental conditions [11]. Evidently, there is a lack of a framework that treats these capabilities as parts of an autonomous computing system. Notably, some of the clearest early attempts to connect them appear not in the academic literature but in patents [13, 17, 24]. In these patents, energy conversion, thermal management and computational demand were considered together more than a decade before comparable integrated developments emerged in research. An important challenge is therefore no longer whether computing can be placed offshore, but how computing and marine infrastructure can operate as one autonomous system.

2. Related Surveys, Patents, and Our Contribution

The literature relevant to autonomous oceanic computing comes from three main sources. They include surveys covering parts of the design space, studies providing the technical foundations on which this survey builds, and patents describing integrated system concepts that have not yet been examined systematically in the review literature. We include surveys whose scope overlaps with at least one key dimension of autonomous oceanic computing, including marine deployment, marine energy supply, thermal management, connectivity, autonomy or positioning. By contrast, we exclude work focused primarily on ocean observation, underwater sensor networks or marine data analytics. Such studies use computing to observe and understand the ocean, whereas our focus is on the ocean as a location, energy source, and operating environment for computing infrastructure.

2.1. Surveys and Primary Studies

No existing survey addresses autonomous oceanic computing platforms directly. The comparison is therefore drawn against four adjacent bodies of work, each of which covers one facet of the paradigm.

1) Surveys on submerged data centers.

Research on submerged data centres provides the closest existing foundation for oceanic computing, but it covers only part of the design space. One review provides a broad overview of submerged-facility architectures, enabling technologies, operational challenges, automation and maintenance requirements [23], while another examines underwater deployment powered by offshore wind, solar and wave energy as a possible route towards zero-carbon artificial intelligence [25]. In both cases, however, the data centre is essentially a sealed facility anchored to the seabed and

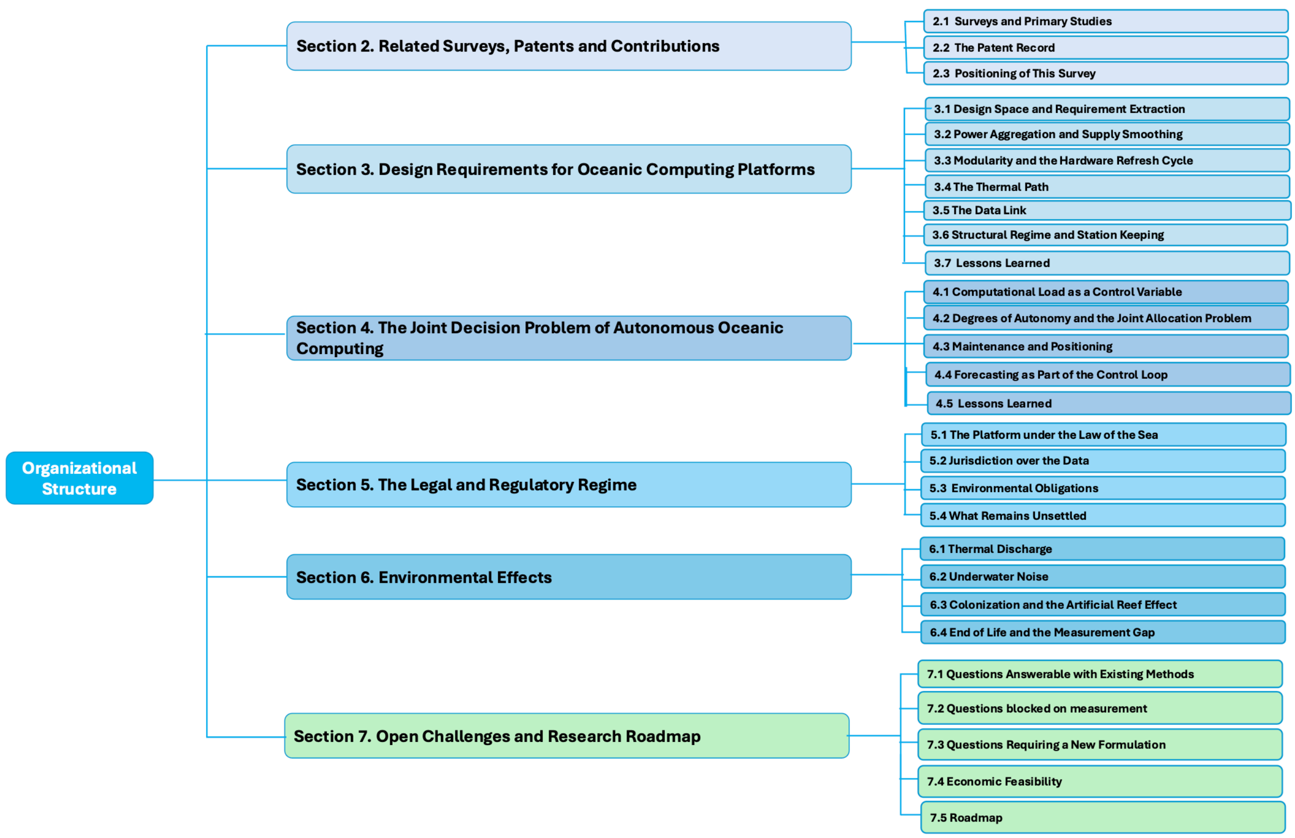

Figure 1: Structure of this survey

connected to shore for power and data. Floating platforms, station-keeping systems and relocatable facilities therefore remain outside their scope [14, 17]. More importantly, neither study considers the fully off-grid operation considered in recent proposals [17], in which a platform must generate its own energy, manage its own connectivity, and potentially change its location. Physical position consequently remains a fixed property of the infrastructure instead of an operational decision variable.

2) Surveys on terrestrial data-center sustainability and thermal management.

Terrestrial data-centre research provides an important baseline for evaluating marine deployment [3, 29]. Existing surveys cover a wide range of approaches for improving sustainability, including computing virtualization, workload consolidation, renewable-energy integration, waste-heat reuse, modular deployment and workload migration across geographically distributed facilities [6]. Other work has organized these techniques around the broader goal of carbon neutrality, treating computing, energy supply, and cooling as parts of the same system [7]. A related body of literature focuses specifically on heat removal, from the energy-saving potential of individual cooling technologies [8] to broader classifications of passive and active cooling strategies [9].

Despite this systems perspective, these studies largely assume a terrestrial operating environment. The data centre remains connected to the grid, physically accessible and fixed at a land-based site. Seawater can be a heat sink. When location is considered, it is usually treated as a long-term siting decision, e.g., where to build a facility, where to place a modular unit, or which site should receive a migrated workload. Thus, the location is not a variable that can change during operation. Recent studies have begun to integrate IT workloads, thermal management, renewable-energy supply and energy storage within a common framework [4, 28]. However, they still cover only part of the oceanic-computing problem because terrestrial facilities can rely on fixed connectivity, routine physical access, and a stationary location. In oceanic computing, these assumptions no longer necessarily hold, and communication, maintenance and position become operational decisions in their own right. Autonomy therefore becomes not simply an additional capability, but a requirement for coordinating computing infrastructure with a changing physical environment.

3) Reviews on maritime autonomous systems and digital twins. Maritime engineering provides many of the methods needed to operate complex assets at sea. Ships and offshore structures are increasingly being developed as self-aware cyber-physical systems, supported by continuous sensing, artificial intelligence, high-speed communication and digital-twin technologies [16]. These capabilities allow structural condition, machinery health, environmental conditions, and operational risks to be monitored throughout the asset life cycle. This approach shifts maintenance from periodic inspection and reactive intervention towards continuous health monitoring and predictive management. The maritime digital-twin literature captures this progression through a sequence of capabilities, from descriptive monitoring and diagnosis to prediction, prescription, and ultimately autonomous operation [15]. More broadly, future marine platforms have been envisioned as mobile power plants, data centres, logistics hubs, and components of larger floating industrial ecosystems [16]. This literature, however, is primarily concerned with the operation of the marine platform itself. It asks how the structure should sense, predict, and respond to its physical environment, but generally treats the computing payload as an external or passive function. Workload demand, service deadlines, data movement, and computational priorities therefore remain outside the platform control loop. For oceanic computing, this separation is no longer sufficient: the physical platform and the computing system must be optimized jointly.

4) Surveys on maritime connectivity and marine energy conversion. Maritime connectivity has been widely studied for vessels and coastal infrastructure, where communication mainly supports navigation, safety, and logistics [21]. An off-grid computing platform creates a different requirement because the link carries computational results whose volume, urgency, and tolerance to delay depend on the workloads executed on board. Communication therefore becomes part of workload management. A similar gap exists on the energy side: reviews of offshore renewable-energy systems characterize the available resource and its conversion into electricity [11]. However, they typically treat the load as an external demand. In oceanic computing, computational demand can instead adapt to, and potentially influence, the energy-conversion process.

2.2. The Patent Record

A substantial body of engineering knowledge on oceanic computing lies outside the peer-reviewed literature. The patent record spans several jurisdictions and, for the systems considered here, extends from 2007 to 2025. We define the corpus by claim scope. Specifically, we include system-level disclosures that treat the energy supply, thermal path, and computing workload as parts of one apparatus. We exclude patents directed only to individual components, such as mooring systems, pressure vessels or wave-energy converters, when they neither incorporate nor serve a computing component or load. We also exclude terrestrial data-centre technologies that were only later proposed for marine use.

Four patent families are particularly important as they describe integrated offshore computing architectures. The earliest, assigned to Google, describes a floating data centre powered by an array of wave-energy machines, cooled through seawater-to-freshwater heat exchangers and built from crane-removable computing modules [13]. IBM later disclosed interconnected floating compute nodes that generate and share power, and that can form a rigid linkage when wave energy exceeds a threshold [30, 31]. Patents associated with Microsoft's Project Natick focus on submerged facilities, including integrated heat exchangers, intrusion detection, and the use of the pressure vessel as an artificial reef [32–34]. A further family assigned to Lone Gull Holdings moves towards greater autonomy: wave-powered computing modules can operate in farms, share power over a common bus [35], and regulate turbine operation by changing the rate of computational work [17, 24, 36]. Comparable system-level disclosures also appear from Fujitsu [37], the Korea Institute of Ocean Science and Technology [38], Keppel Data Centres [39, 40], Ciel et Terre [41], Nautilus True [42] and Shenzhen Hailanyun [43]. Most of these later systems are externally powered and mounted on the seabed or on floating platforms, whereas self-powered concepts appeared as early as 2007.

The patent record is valuable because it captures integrated design requirements that the academic literature has not yet fully characterized. A granted patent is tied to a filing date and has been examined for novelty, non-obviousness, and sufficiency of disclosure [44, 45]. Also, its specification must describe the invention in sufficient detail for a skilled practitioner to implement it. We therefore treat granted patents as evidence of engineering concepts that had reached a defined level of technical formulation. Note, however, that pending applications are treated as statements of intent, and abandoned applications as evidence of what was possibly conceived. Patent analysis has already been used to trace the evolution of data-centre cooling technologies [46], and we apply it here on the same methodological basis. The resulting requirements are developed in Section 3 and summarized in Table 1. Importantly, no reviewed work covers the patent record as a whole, and although positioning and energy–compute coupling both appear in the literature, they are not treated together. This separation is one of the gaps that motivates the present survey.

2.3. Positioning of This Survey

Table 1 positions this survey at the intersection of several research communities that have largely developed independently. Existing studies independently examine submerged data centers, terrestrial data-center sustainability, marine renewable energy, maritime connectivity, or autonomous offshore platforms. However, they do not treat these elements as components of a single computing infrastructure. This survey addresses that gap through a unified view of oceanic computing. Specifically, energy harvesting, workload management, thermal control, communication, and maintenance are coordinated as domains of one cyber-physical decision problem. Where applicable, positioning is coordinated as part of the same problem. These domains share limited energy, storage, communication, and maintenance resources and interact across timescales ranging from the wave period to the hardware refresh cycle. Consequently, a decision in one domain changes the feasible decisions in the others.

Methodologically, the survey combines three complementary sources of evidence: peer-reviewed research, documented deployments, and system-level patents. The patent record provides dated, technically specific disclosures of integrated design requirements that often appeared before they were formalized in the academic literature. The progression is revealing. A wave-powered computing array was disclosed in 2007 [13]; interconnected floating nodes capable of generating and sharing power followed in 2014 [30]; self-powered computing buoys connected through a common power bus were filed in 2017 [24]; and a later filing linked turbine resistance directly to the rate of computational work [17]. The last development is particularly significant because it turns computational demand into a control variable of the energy-conversion system. For each requirement identified in the patent record, we then assess what has been demonstrated, what remains conceptual, and what is still uncharacterized at system level.

The scope of the survey also extends to constraints that are usually treated outside the core data-center literature. Acoustic injection attacks against underwater hard-disk storage [47, 48], together with hydrophone-based acoustic localization for mobile robots [26], connect enclosure design, physical security, sensing, and maintenance. Platform mobility creates another coupling. Movement can change the applicable maritime regime [49]. At the same time, data-protection obligations may follow the operator and processed data regardless of location [50, 51]. Environmental and economic evidence, however, remains limited. Our literature search found no published field measurements quantifying the marine effects of an operating offshore computing facility. Additionally, the public record lacks a complete economic assessment of the autonomous configurations examined in this survey.

Accordingly, the objective of this survey is to establish autonomous oceanic computing as a coherent object of study and to identify the integrated decision problem that defines it. To the best of our literature search, no prior review jointly treats marine renewable energy, autonomous platform control, and computing infrastructure, despite emerging systems that sit precisely at their intersection [52, 53]. We therefore consolidate the engineering requirements, their coupled operational decisions, and the legal, environmental, and economic conditions shaping the paradigm.

The main contributions of this survey are summarized as follows.

- We provide a unified design space for offshore computing facilities along four construction-time axes: siting, structural regime, energy coupling, and data link. This framework places submerged pods, moored platforms, converted vessels, and self-propelled nodes within a single taxonomy rather than treating them as unrelated branches.
- We reconstruct the evolution of integrated oceanic-computing architectures through the patent record and extract the design requirements that it establishes. Particular attention is given to the progression from the co-location of computing and marine energy generation to closed-loop energy–compute control, in which computational demand influences energy conversion. We then assess each requirement against the corresponding peer-reviewed engineering literature.
- We characterize operational autonomy across six domains: energy harvesting, workload management, thermal control, communication, maintenance, and positioning. For each domain, we review the available methods and assess the degree of autonomy reported for existing systems. We further show that autonomy achieved independently within individual domains does not necessarily compose into autonomy of the complete platform.
- We examine the conditions external to the core technology, including platform classification under maritime law, jurisdiction over processed data, environmental effects, physical-layer security, and economic feasibility. We then organize the resulting open problems into a research roadmap according to their methodological and measurement prerequisites.

In summary, these contributions distinguish oceanic computing from a conventional data center relocated offshore. Once a platform supplies its own energy and, in some configurations, controls its own position, its physical and computational subsystems cease to be separable. Design choices in one domain become operating constraints in the others, and the conventional boundary between infrastructure design and infrastructure operation begins to dissolve. The survey is organized as illustrated in Figure 1.

3. Design Requirements for Oceanic Computing Platforms

3.1. Design Space and Requirement Extraction

Terminology for marine computing infrastructure remains inconsistent. The terms underwater, subsea, floating, and offshore may refer to water-column placement, structural form, or geographic location, and similar facilities are consequently classified differently across prior studies [1, 23, 25, 26]. We use oceanic computing as the umbrella term for this broader class of infrastructure. As shown in Figure 2, individual facilities are classified along four logically independent construction-time axes: i) siting, from coastal and nearshore waters to the high seas [54]; ii) structural regime, from sealed seabed pods to self-propelled nodes; iii) energy coupling, from a terrestrial-grid connection to onboard energy harvesting [24]; and iv) data link, from subsea fibre to satellite communication [21]. Every documented or disclosed facility occupies one position on each axis. Examples range from the sealed seabed deployments represented by Project Natick and the Lingang facility [14, 55] to self-propelled, wave-powered nodes using satellite backhaul [52, 53]. Reading the axes from left to right therefore

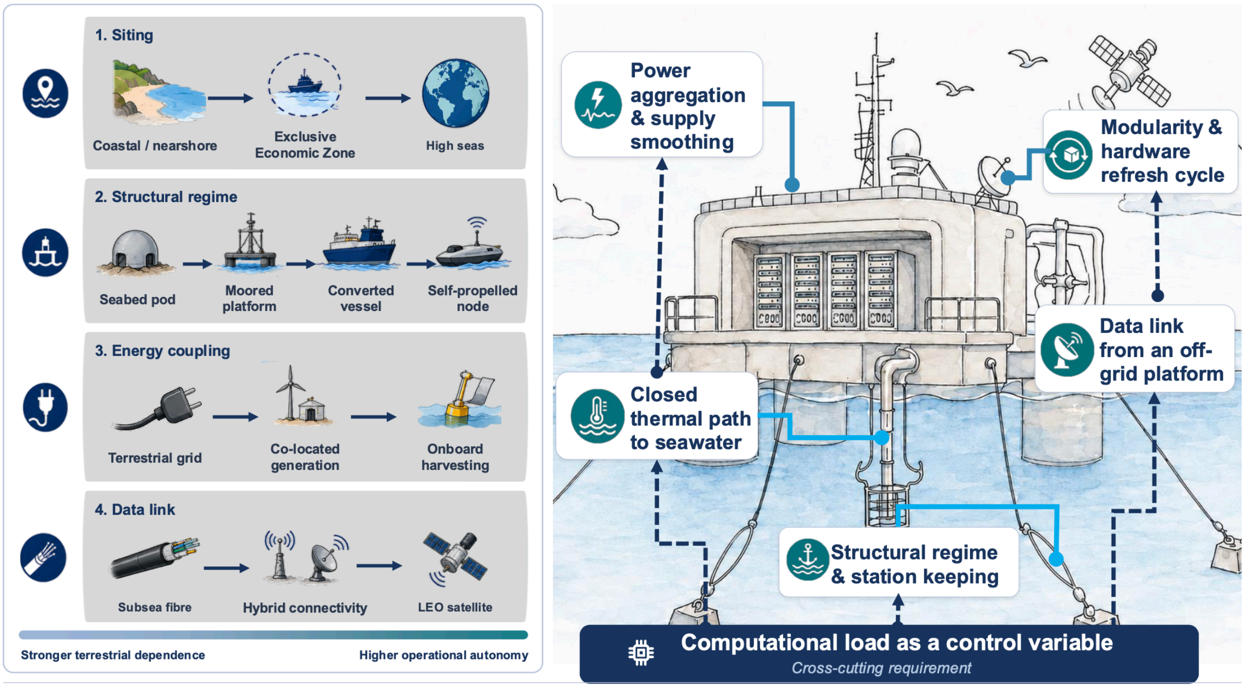

Figure 2: Design space and engineering requirements of oceanic computing platforms. Construction-time choices progressively reduce dependence on terrestrial infrastructure and introduce corresponding system-level requirements. Computational load acts as a cross-cutting control variable across the entire design space.

indicates increasing detachment from terrestrial infrastructure. The engineering evidence associated with this design space is fragmented across several research communities. Peer-reviewed studies address individual subsystems, including data-center cooling and thermal management [8, 9], marine energy conversion and power smoothing [56, 57], and maritime platform autonomy [15]. By contrast, several architectures integrating the energy supply, thermal path, computing payload and marine structure first appeared in patent disclosures [13, 24, 30–36]. Patent-based analysis has previously been used to trace the evolution of data-center cooling technologies and to identify shifts in engineering priorities before they become prominent in the publication record [46]. We therefore use the patent record to identify system-level requirements, regardless of their technical or commercial viability. Peer-reviewed engineering studies are then used to assess whether each requirement can be met, under what conditions, and with which unresolved limitations.

Section 2.2 has already reconstructed the relevant patent families and their chronology. The purpose here is therefore to translate the disclosures into engineering requirements. The onboard-energy-harvesting patents describe floating or self-powered computing arrangements in which energy generation, thermal management, and the computing load form parts of one apparatus (Figure 3) [13, 17, 24, 30, 31, 35, 36]. The submerged-data-center patents contribute requirements associated with sealed pressure vessels, heat transfer through or across the enclosure, and combined power and data connections [32]-[34].

Sections 3.2-3.6 examine the first five requirements by pairing each patent-derived requirement with the corresponding engineering evidence and identifying what remains unresolved for an autonomous off-grid platform. The sixth requirement—the use of computational load as a control variable of the energy-conversion system [17]—cuts across the other five and is therefore developed in Section 4.1 as the foundation of the joint decision problem.

3.2. Power Aggregation and Supply Smoothing

Wave power is commonly represented as a narrow-band stochastic process [58]. At a single converter, electrical output oscillates with the wave period, while its envelope changes with the sea state. The resulting peak-to-mean power ratio can be large, requiring the conversion system to be rated above its average output [58]. A wave-powered computing platform must therefore manage both the average energy balance and short-term supply-demand mismatch.

The patent record identifies two complementary supply-side mechanisms. Google's 2007 filing replaces a single large converter with multiple smaller wave-motion machines arranged in an electrically interconnected grid, i.e., an array [13]. The self-powered buoy family uses a different form of aggregation: mutually inter-tethered converters exchange power over a common bus [24]. Because spatially separated units experience different wave phases, their outputs can

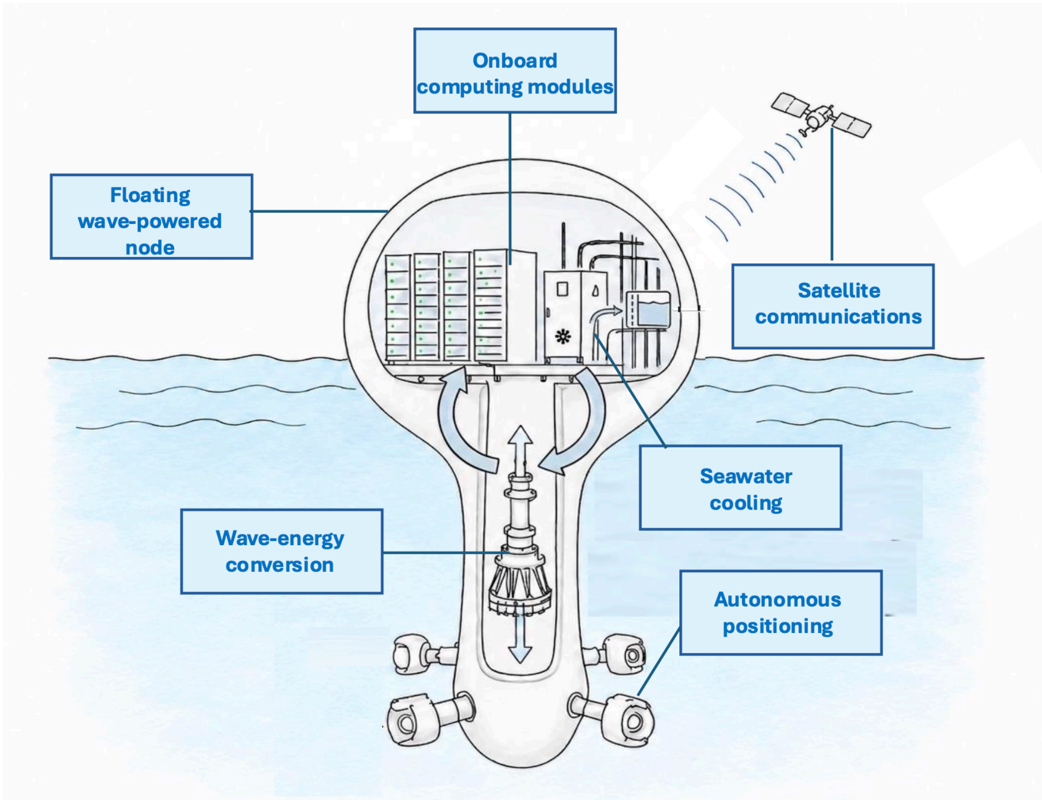

Figure 3: Example of a floating, wave-powered oceanic computing node. Wave energy is converted on board to power the computing modules, while the computational load can in turn influence the operating point of the energy converter [17]. Seawater provides cooling [8], satellite links provide connectivity [21], and autonomous positioning supports station keeping and relocation [20, 24].

partially compensate for one another, reducing aggregate power fluctuations and onboard storage requirements.

The common-bus arrangement can also simplify the electrical conversion chain. A wave-energy generator generally produces variable-frequency alternating current, whereas computing hardware requires regulated direct current. In a conventional architecture, the converter output is rectified, converted back to alternating current for distribution, and rectified again at the rack. A shared direct-current bus removes the intermediate alternating-current stage, reducing conversion losses [24]. Thus, spatial aggregation smooths generation, whereas bus integration enables power sharing and shortens the conversion chain.

Studies of wave-energy arrays quantify how smoothing depends on converter count and spatial arrangement. Aggregate variability decreases as converters are added, with the largest gains for small arrays and diminishing returns as the array grows [56]. For large arrays, power variance has been reported to decrease by approximately a factor of five [57]. Geometry also matters: the normalized power variance of rectangular arrays was reported to be approximately twice that of circular arrays with the same number of converters [57]. Converter count and spatial layout therefore determine the variability presented to the computing platform. This is particularly relevant to the architecture in [35], which comprises eight self-propelled, energy-harvesting computing nodes operating as a coordinated cluster. It therefore falls within the small-array regime in which the marginal smoothing benefit is most pronounced [56]. By distributing the nodes spatially and combining their generated power, the cluster can reduce the variability experienced by the computing system.

Aggregation nevertheless has two important limitations. First, spatial smoothing does not provide the regulated direct-current voltage required by computing hardware. A power-electronic interface remains necessary, and spatial distribution is consequently regarded as one smoothing measure [59]. Second, both array smoothing and common-bus power sharing require multiple converters. They therefore do not apply to standalone systems such as the self-powered computing buoy [24] and the initial Panthalassa node [52]. A standalone platform must instead accommodate the full variability and peak-to-mean ratio of its own converter through storage, curtailment, flexible demand, or a combination of these mechanisms.

Energy storage provides one response. Hybrid architectures commonly use supercapacitors to absorb rapid power fluctuations and batteries to manage slower variations [60, 61]. In one dispatching study, an optimized hybrid configuration achieved a cost of 2.6 cents per kWh, compared with 8.4

cents per kWh for battery-only storage [62]. For an oceanic computing platform, however, storage design also has important maintenance implications. Battery degradation may require vessel deployment, personnel transfer, and hardware replacement, causing direct costs and possible service interruption. Larger storage capacity can reduce curtailment, but increases platform mass, capital cost, and equipment requirements.

Flexible computing demand provides a complementary mechanism. The rate of computational load can be adjusted so that electrical demand follows the available energy. In the architecture disclosed in [17], load modulation changes the electrical load seen by the generator. Unlike aggregation and power sharing, load modulation remains available to a standalone node. For a fleet, it can operate together with supply-side smoothing: power is aggregated and shared across converters or computing nodes [35], while computational demand is adjusted over time [17]. Spatial aggregation smooths supply, whereas load modulation reshapes demand.

To the best of our review, these mechanisms have not been designed jointly. Wave-energy studies optimize converter count and array geometry while treating the consumer as a fixed load or a grid connection [56, 57]. Computing studies adapt workload demand while treating renewable generation as an exogenous process. A joint formulation would instead select array size and geometry, storage capacity, and workload flexibility to meet a target on delivered computational work while quantifying their trade-offs.

Power aggregation therefore becomes an autonomous energy-management problem. Unlike a terrestrial data center, an open-ocean platform cannot rely on electricity prices or grid carbon intensity as external dispatch signals. It must instead operate according to its system state, workload commitments, stored energy, and forecast wave input. At each decision interval, it must determine how much power to allocate to computing and auxiliary systems, how much to store, which workloads to execute or defer, and when to curtail energy capture.

Existing online formulations provide partial foundations for these decisions. Some coordinate renewable generation and storage across grid-connected data centers [63], while others address fully energy-harvesting edge systems using server-frequency scaling [63]. In both cases, however, energy supply is treated as exogenous and computing adapts to it. Oceanic computing introduces a bidirectional relationship: energy constrains computational load, while load changes can alter generator torque and the operating point of the conversion system [17]. This closed-loop energy–compute coupling has not yet been characterized and is developed further in Section 4.1.

3.3. Modularity and the Hardware Refresh Cycle

A marine computing facility combines assets with fundamentally different lifetimes. The offshore structure is typically designed for decades of service, whereas computing hardware, especially Graphics Processing Units (GPUs) and Neural Processing Units (NPUs), may become obsolete within a few years [14]. Whether computing equipment can be replaced independently of the structure is therefore a construction-time decision. The 2007 floating-data-center disclosure enables this through crane-removable computing modules [13]. Sealed submerged designs instead prioritize enclosure integrity and component protection, tying hardware refresh to recovery of the entire pod.

Project Natick illustrates the reliability benefit of the sealed approach. The vessel was filled with a controlled nitrogen atmosphere and operated without routine human access, reducing oxygen, humidity, and human handling as sources of failure. Over its two-year deployment, the reported component failure rate was much less than that in a comparable land-based facility [14]. The sealed configuration, however, eliminates reconfigurability during deployment. The balance among computing, storage, and accelerator resources remains fixed, and hardware upgrades require pod recovery. Thus, the pod can operate unattended but cannot perform autonomous physical maintenance or hardware reconfiguration. The evidence also spans two years, whereas the structure may remain serviceable for much longer.

Floating and surface-accessible configurations preserve greater optionality. Individual modules can, in principle, be removed, replaced, or upgraded while the supporting structure remains in service [13, 24]. Computing and storage capacity can be rebalanced, and newer accelerators introduced without replacing the complete platform. These benefits come with marine-exposed interfaces, including hull penetrations, connectors, lifting points, and sealing surfaces near the splash zone. Surface access therefore enables intervention, but not continuously. The hardware-refresh interval depends on module design, vessel capability, and suitable weather windows.

The offshore-wind literature provides a useful basis for quantifying this constraint because access depends mainly on the vessel and transfer system [63]. Conventional crew-transfer vessels are typically limited to significant wave heights of approximately 1.5–1.8 m. At the North Sea site studied in [63], this corresponded to an annual access probability of about 0.65, leaving such vessels unavailable for roughly 30–35% of the year. Service-operation vessels, operable up to about 3.2 m, raised the probability to approximately 0.84, while helicopter access reached about 0.91 subject to wind and visibility constraints [63]. Annual averages, however, mask seasonal concentration: runs of up to five consecutive days without an operational window were identified in January [63].

Module replacement is more restrictive than personnel transfer because it requires offshore lifting. A suspended module moves relative to both the vessel and the platform, and these motions must remain within the limits of the crane, rigging, and receiving structure. Offshore lifts are therefore permitted only below specified wave-height thresholds. Under DNV-ST-N001, these limits are further reduced by an alpha factor accounting for forecast uncertainty [64]. Serviceability therefore depends on vessel access and lifting capacity, platform constraints, and a forecast window that permits safe completion.

Another trade-off is between reliability and optionality. A sealed pod reduces routine intervention but fixes the configuration for the deployment interval. An accessible modular platform allows repair, rebalancing, and payload refresh, but exposes interfaces and makes maintenance sea-state dependent. The patent record includes both approaches [13, 24, 34], but does not quantify which is preferable across failure and obsolescence rates, vessel costs, and access probabilities. Assessment should also consider whether predicted failures are handled by workload migration, onboard cold spares, or intervention during the next feasible weather window.

Marine robotics provides partial foundations for reducing dependence on human access. Resident subsea robotic systems have been developed for inspection and maintenance [65]. Autonomous intervention has been demonstrated for valve turning and plug connection [66]. These results show that robots can perform structured marine tasks, although autonomous replacement of computing modules has not yet been demonstrated. The reported systems also operate around seabed infrastructure supported by shore-based power, communication, or docking, whereas an off-grid platform may need to provide these functions itself.

Self-maintaining computing systems offer a complementary precedent. Software-controlled robotic mechanisms can inspect, service, or repair hardware [67], while mobile manipulators can replace data-center components [68]. Offshore deployment would require onboard robotic support, payload stabilization, and operation under changing sea conditions. Current evidence supports inspection, diagnosis, and selected manipulation, but not autonomous computing-module replacement at sea. A practical architecture would therefore combine reliability, redundancy, predictive failure detection, resident inspection, and weather-aware intervention planning.

3.4. The Thermal Path

Seawater is both the main reason marine data centers are attractive and the main constraint on how they can be cooled. For the same volume and temperature rise, water can absorb roughly 1,000 times as much heat as air. This high volumetric heat capacity allows an offshore facility to release server heat without mechanical chillers. Yet seawater is also an aggressive electrolyte. Its chloride content promotes pitting and crevice corrosion, so every surface that comes into contact with it must be made from a corrosion-resistant alloy. Titanium is the benchmark material for heat-exchanger tubing [69].

Accordingly, every design examined here isolates the computing equipment within a closed cooling loop and transfers heat to the environment through a defined external interface. The designs differ mainly in where this interface is placed: it may be a replaceable plate heat exchanger, or it may be integrated into the hull itself. In the 2007 disclosure, fresh water circulates through the internal loop and carries heat away from the electronics. The heat then passes to seawater through seawater-to-freshwater heat exchangers [13]. Seawater never reaches the computing equipment, and the internal coolant can be selected to match the requirements of the electronics.

Microsoft's submerged design takes a different approach, using integrated and external heat exchangers on the vessel hull [32]. In this configuration, the hull forms part of the thermal path between the servers and the surrounding sea. This shortens the distance over which heat must travel and avoids the pumping power required by a separate seawater circuit. The trade-off is that the heat-transfer surface is also the surface on which marine organisms accumulate.

This process, known as biofouling, develops in stages. Organic molecules first form a conditioning film on the immersed surface. Bacteria then attach to this film and secrete a polymeric matrix, creating a biofilm on which larger organisms—including algae, barnacles and tubeworms—can settle [22]. As this layer grows, the initial benefit of a short thermal path declines. The fouling layer acts as additional insulation, adding thermal resistance to that of the hull or exchanger wall. The effect can be substantial: a biofilm only 1 mm thick has been reported to halve heat-transfer efficiency, a greater loss than that caused by most forms of mineral fouling [69].

Heat exchangers are therefore designed with additional capacity from the outset. This "fouling allowance" preserves the required cooling performance as deposits build up between cleanings [69]. The approach, however, assumes that the surface can eventually be reached and cleaned. For a sealed offshore structure, cleaning requires vessel intervention as part of a maintenance cycle. Its fouling allowance must therefore cover not merely the interval between routine cleanings, but the entire deployment period [22].

A self-powered computing platform provides a third configuration. Because it sits at the sea surface, it can exploit both the surrounding water and ocean winds to reduce the cost and complexity of cooling its computing circuits [24]. A submerged pod cannot use this air-side route. Air cooling is also thermodynamically less effective per unit volume of fluid moved. The volumetric heat capacity of air is about three orders of magnitude lower than that of water. Consequently, much larger airflow rates are needed to remove the same amount of heat [70]. Its advantage is that it does not depend on the seawater interface. It therefore degrades for different reasons and under different environmental conditions.

A surface platform equipped with both water-side and air-side cooling can optimize them. This flexibility is important because water-side performance is not constant. Seawater temperature changes with season and location, and marine heat waves reduce the temperature difference that drives heat from the facility into the sea. The resulting loss of cooling margin has even been proposed as a reason to relocate a facility [20]. Biofouling causes a second, slower decline by progressively insulating the submerged exchange surface. Air-side cooling is affected by atmospheric conditions instead, but its performance can recover when those conditions improve.

All three configurations therefore use a closed internal loop that terminates at an external heat-transfer interface. The central design choice is where that interface should be placed and which environment it should use. Each option brings a distinct mode of degradation and a corresponding maintenance requirement. A submerged exchange surface accumulates biofouling and cannot be reached without vessel intervention. A separate plate exchanger remains exposed to corrosion, but it can be replaced. An air-side path depends on ambient conditions, yet can recover as those conditions change.

Autonomy becomes important when a platform has more than one cooling path. By monitoring seawater temperature, fouling, atmospheric conditions and the performance of each route, the platform can shift the cooling load between water-side and air-side systems as conditions change. If neither path can remove enough heat, it can also reduce the computational load and therefore the heat generated by the servers. Thermal management thus becomes part of the same autonomous control problem as energy management.

3.5. The Data Link

The reviewed designs differ fundamentally in how data is returned to shore. Submerged systems assume a cable: the pressure vessel is connected to land for both power and data [14, 32]. The autonomous branch removes this physical connection and instead uses a low-Earth-orbit (LEO) satellite link [52]. This change enables operation on the high seas, but it does more than replace one transmission medium with another. A cable-connected facility remains physically tied to shore, receiving electricity and exchanging data through a fixed link. An off-grid platform generates its own electricity, so the shore connection carries information rather than power, mainly the results of computation.

A LEO satellite link is not equivalent to fibre. Measurements from users of the largest operational constellation show that both throughput and latency vary substantially over time [71]. Performance also degrades periodically, on a roughly 15-second cycle, as terminals switch from one satellite to another. Weather introduces further variation, and service outages remain frequent [71]. These effects are most severe where no ground gateway is nearby, especially for a platform far offshore. A terminal in the middle of the Indian Ocean, which depended on inter-satellite links, showed much greater delay, jitter and packet loss than a terminal close to a gateway. Round-trip times exceeded 500 ms, whereas a more direct route would have remained below 100 ms [72]. In another study, a maritime terminal off the Brazilian coast was separated from shore by a distance corresponding to only about 14 ms of propagation delay, yet its average round-trip time was 96 ms. Sustained increases in latency occurred during 16% of the observation period [73].

These measurements have two consequences that are rarely made explicit in the offshore-computing literature. The first is that the data link determines which workloads are practical. It limits how much information can enter or leave the platform in a given time. A task whose input must be transferred from shore is therefore constrained by the backhaul—the connection between the platform and the wider network. A task that generates its own input on board, or returns only a small result, is much less affected. Off-grid operation consequently favours workloads that perform many computations for each byte transferred, or equivalently, those with a high ratio of computation to communication. The link is not merely a transport detail added after workload selection; it becomes one of the criteria for admitting a workload in the first place.

The second consequence concerns scheduling. Link capacity varies over timescales ranging from seconds to hours, in response to weather and satellite geometry [71]. At the same time, a wave-powered platform may already be changing its computational demand to follow a variable energy supply [17]. These two constraints are independent: strong wave conditions do not necessarily coincide with good satellite connectivity. The platform may therefore have abundant energy for computing but no reliable route for returning the results. Outputs must be stored on board until a suitable transmission window appears, making storage a scheduling resource rather than simply a repository. The scheduler must decide not only what to compute, but also which results to retain, for how long, and at what level of compression. The choice of connection also changes the physical attack surface, because a fixed cable and a wireless terminal expose the platform in different ways.

Free-space optical communication provides a third option for high-capacity satellite links [74, 75]. Instead of using radio waves, it carries data on a narrow laser beam through the atmosphere. It requires no physical connection to shore and is therefore compatible with a mobile platform. Optical links can provide data rates orders of magnitude higher than radio-frequency links, do not require the same spectrum licensing, and are intrinsically harder to intercept because the beam is tightly confined [76]. The same narrow beam, however, creates the central engineering challenge. The terminal must first acquire the satellite, track its motion and keep the laser accurately pointed at it. On a moving marine platform, acquisition, tracking and pointing can become the limiting subsystem [76].

Experiments show that such tracking is feasible at sea. A shipborne acquisition, tracking and pointing system has acquired and followed LEO satellites from a vessel, reporting a pointing accuracy below 180 μrad when referenced to stars and a tracking accuracy of about 20 μrad against satellites [77]. Moving-platform optical links have also reached gigabit-per-second rates between an aircraft terminal and a satellite in geostationary orbit under cloud-affected and changing atmospheric conditions [78]. These demonstrations indicate that the essential requirements are an accurate reference for the platform's orientation and a stabilized terminal. A platform that already performs station keeping and attitude control can, in principle, provide both.

One asymmetry nevertheless works against the offshore case. In satellite optical communication, the uplink—from the platform towards space—is generally harder than the downlink. An uplink beam encounters the densest part of the atmosphere immediately after transmission. Spatial and temporal variations in the refractive index can then distort the wavefront and introduce pointing instability. A downlink, by contrast, is affected more strongly by geometric beam spreading [79]. An offshore computing platform mainly exports results, so most of its traffic is in the more difficult uplink direction. The challenge is therefore not simply to transfer optical technology from a satellite downlink to a platform at sea, but to operate it in the direction in which atmospheric propagation is least favourable.

Two results from space optical communication help to define how this challenge might be managed. The first concerns availability. An optical beam cannot pass through cloud, and no improvement in pointing can remove that limitation. The conventional response is site diversity: several ground terminals are placed far enough apart that their cloud conditions are approximately independent, and the system selects whichever station has a clear path [75]. Studies of high-capacity optical feeder links therefore treat the number and geographical separation of ground terminals as the main design variables determining overall availability [75].

An offshore platform can partly extend this idea as the transmitting terminal is itself mobile. A fixed transmitter can improve availability mainly by adding receiving sites. A relocatable ocean platform can also move towards a clearer atmospheric path or more favourable link geometry. Diversity can therefore be created both by multiplying receivers and by repositioning the source. In this limited sense, mobility gives the platform a form of site diversity that a fixed installation must obtain through additional ground infrastructure.

The second result concerns how to use a link that is available only intermittently. The TBIRD mission demonstrated a 200-Gbps optical downlink from a 6U CubeSat in LEO. It transferred as much as 4.8 TB without error during a single five-minute pass, using a 2-TB onboard storage buffer [74]. Two features of the architecture are particularly relevant. First, the transfer is organized from buffer to buffer: data accumulate on the spacecraft and are sent in a high-rate burst when the orbital geometry permits. Second, reliability is provided at the transport layer through an automatic repeat request (ARQ) protocol. Acknowledgements are returned on a low-rate optical uplink. This arrangement supports error-free delivery under atmospheric fading and allows a transfer to complete even when many acknowledgement messages are lost [74].

An offshore computing platform faces the same basic pattern: a transmitter accumulates data away from shore and a receiver waits on land. The main difference is that the transmitter sits on the sea surface rather than in orbit. The implications for platform design are direct. Onboard storage must hold completed results until the next usable communication window, so the required capacity depends on the longest likely waiting time before transmission, not only on the working set of the computation. During that interval, the scheduler must decide which tasks to execute and which data to preserve, compress or discard. The demonstrated ability to move terabytes in a single pass makes this regime plausible even for workloads that generate large outputs. It therefore broadens the potential workload set beyond tasks that return only compact results.

Transferring this architecture from space to a marine platform introduces additional challenges. TBIRD communicates between a spacecraft whose attitude is controlled independently and a fixed ground station on land. An ocean platform must keep its beam aligned despite wave-induced translation and rotation, propagation through the marine boundary layer—the turbulent region of air just above the sea—and operation in the uplink direction, where atmospheric distortion is strongest. Link availability therefore depends on how accurately the platform can maintain its angular orientation in a given sea state.

This dependence has been quantified for ship-to-shore optical links [80]. One analysis represented the vessel's six oscillatory motions using its response amplitude operator, a standard model of how a vessel moves in response to waves. It then evaluated geometric spreading, atmospheric attenuation, rapid intensity fluctuations known as scintillation, and pointing misalignment together [80]. Although that study considered a beam directed towards shore rather than towards a satellite on the high seas, its central dependence remains: whatever the destination, platform motion determines pointing stability.

This creates a direct coupling between communication and platform control. With a stabilized optical terminal, backhaul capacity depends on accurate pointing. Accurate pointing, in turn, requires the platform to control its position and orientation. Energy spent on station keeping or attitude correction is then unavailable for computation, while a change in heading or location may improve a communication opportunity. Communication therefore becomes part of the same control problem as positioning and energy management, rather than a separate networking function.

Autonomy in this domain concerns three questions: which link to use, when to transmit, and what to transmit [21, 81]. The platform cannot depend on continuous instructions from shore because the opportunities for communication may themselves be brief. Satellite passes can be predicted from orbital information, but the platform's orientation during a pass depends on the sea state. The same wave forecast used to manage energy conversion can also help predict vessel motion and, therefore, pointing conditions [82, 83].

When several satellites are visible, the platform must decide which one to track. A self-propelled platform may also adjust its heading or orientation to support the selected pass. Because acquisition and tracking take time, the terminal must begin aligning before the communication window opens. The platform must also rank the data waiting in its buffer. It has to decide which results should be transmitted first, which can be compressed, and which may be discarded when storage or contact time is limited.

If both optical and radio-frequency links are available, the platform must choose between them using forecasts of weather, sea state and link quality. The optical link offers much greater capacity, but requires a clear atmospheric path and precise pointing. The radio-frequency link is slower, but is more tolerant of cloud and platform motion. To our knowledge, existing work has not yet used the same wave forecast both to control energy harvesting and to anticipate communication pointing. Closing this gap would allow energy production, platform motion and data return to be planned through a common forecast rather than as three separate functions.

3.6. Structural Regime and Station Keeping

The designs reviewed here treat position in three different ways. A submerged pod is anchored, so its location is fixed when it is installed [34]. A farm of computing buoys is moored, with the individual units also tethered to one another. Their movement is therefore constrained both relative to the seabed and relative to neighbouring units [24]. A self-propelled node has no permanent connection to the seabed. Instead, it uses thrust to hold its position or move to a new one [52]. These arrangements create different structural loads and, more importantly for an off-grid computing platform, very different demands on its limited energy supply.

A platform can maintain position passively through anchors and mooring lines, or actively through thrusters. The energy difference is substantial. Under the same wind, wave and current conditions, one comparative study found that anchor-based positioning consumed approximately 24% less energy than a thruster-based dynamic-positioning system. The lower energy use produced a proportional reduction in exhaust emissions, including sulfur oxides and carbon dioxide [84].

Hybrid systems can recover part of this energy advantage while retaining accurate position control. In thruster-assisted mooring, the mooring lines absorb the relatively steady forces imposed by currents and wind, while the thrusters correct only the slow drifting motion that the mooring system cannot fully suppress. This division of work reduces the energy required to hold position [85]. Mooring-assisted dynamic positioning has similarly been shown to improve positioning accuracy while using less thruster power than either mooring or dynamic positioning alone [86]. The same combination has also been evaluated for floating offshore wind platforms [87].

These comparisons are usually expressed in terms of fuel use or operating cost because conventional vessels obtain energy from fuel tanks or shore connections. For an off-grid computing platform, the trade-off is more direct. Every kilowatt used for propulsion is a kilowatt that cannot be used for computation. Both demands draw from the same harvested supply that the platform or array was designed to smooth [24]. Station keeping therefore becomes part of the same resource-allocation problem as cooling and communication. Its cost is also not constant: stronger waves, winds and currents require greater thrust and leave less energy available for computing.

On the high seas, the lower-energy option may not be available. Mooring requires both a seabed within reach and a jurisdiction in which anchoring is permitted. These conditions may not hold for a platform operating in international waters above the deep ocean. Yet waves, winds and currents continue to displace the structure. A moored platform can resist these forces without continuously consuming propulsion energy. A free-floating platform must run its thrusters whenever it needs to remain in place.

This distinction reverses the usual assumption about infrastructure. Holding position is the energy-intensive state, whereas drifting is the low-energy state. A conventional moored asset treats displacement as a fault that must be corrected. A free-floating computing platform can instead treat drift as its normal mode of operation. It expends propulsion energy only when occupying a particular location provides enough value to justify the cost. That value may come from access to a better energy resource, more favourable geometry for pointing a communication terminal, or proximity to a maintenance base.

A platform may also change its structural state in response to extreme conditions. Offshore aquaculture platforms designed for typhoon-exposed waters use two operating drafts. In normal conditions, the structure remains afloat. In a survival configuration, it is submerged and allowed to rest on the seabed, taking advantage of the reduction in wave effects with depth [88]. This strategy is possible only where the seabed is within reach and therefore does not apply directly to deep-ocean platforms. It nevertheless establishes an important principle: the structural regime shown in Figure 3 need not always be fixed at the time of construction. In some environments, it can become an operating decision.

Autonomy in positioning therefore concerns a basic choice: whether to hold station or allow the platform to drift. The platform must estimate the thrust required to maintain position under the current sea state and compare its energy cost with the value of remaining there [89, 90]. When the expected benefit is lower than the cost, the rational decision is to stop resisting the environment and drift.

3.7. Lessons Learned

This section identifies five design requirements from the patent record. An oceanic computing platform must aggregate a variable energy supply; allow short-lived computing hardware to be refreshed independently of the longer-lived marine structure; transfer heat safely to the sea; return results over an intermittent link; and hold position when necessary. Each requirement is supported by an established engineering literature, but that literature was developed for different systems and objectives. Wave farms feed electricity into a grid. Offshore wind turbines are maintained by dedicated vessel fleets. Ships use fuel to hold position. Satellite terminals generally operate from fixed ground stations. In these

Table 2
Documented offshore computing deployments and disclosed concepts, positioned on the four design axes of Figure 2. Values marked * are declared by the operator or manufacturer.

System	Siting	Structural regime	Energy coupling	Data link	Scale and status
Project Natick, Phase 2 [14]	Nearshore	Sealed seabed pod	Shore grid via cable	Subsea cable	864 servers, ≈36 m depth; 2018–2020; programme closed 2024
Highlander, Hainan [54]	Coastal, EEZ	Sealed seabed pod	Shore grid via cable	Subsea cable	≈1400 t module at 35 m*; first commercial operation
Lingang, Shanghai [55]	Coastal, EEZ	Sealed seabed pod	Co-located offshore wind	Subsea cable	24 MW declared, from a 2.3 MW first phase*; ≈2000 servers*; commercial operation 2026
Subsea Cloud [54]	Coastal	Sealed pod, dielectric immersion	Shore grid via cable	Subsea cable	Commercial service; projects announced*
Nautilus Data Technologies [91]	Coastal, inshore	Moored barge	Shore grid	Terrestrial fibre	Waterborne facility in operation since 2021*
Converted vessels (MOL, Hitachi) [92]	Coastal	Converted vessel	Shore grid or onboard generation	Not disclosed	Memorandum of understanding 2026; operation targeted 2027*
AO60DC [93]	Offshore, EEZ	Moored floating wind platform	Co-located generation, onboard storage	Not disclosed	15–18 MW turbine with 10–12 MW of compute*
Panthalassa [52, 53]	High seas	Self-propelled node	Onboard wave harvesting	LEO satellite	≈85 m node*; pilot deployment preceding commercial operation

systems, the supporting arrangements are largely set at the design stage. An autonomous ocean platform must instead meet the same needs continuously, under conditions that may change from one wave period to the next. Existing methods remain useful, but their original assumptions and objectives do not transfer directly.

The five requirements are tightly coupled. A platform has one hull, one energy budget and only a limited number of suitable weather windows. Energy used for thrust is unavailable for computation. Increasing energy-storage capacity can smooth the power supply, but it also adds mass that the structure must carry. Maintenance can take place only when sea conditions permit safe access. A decision in one area therefore changes what is possible in the others; none of the requirements can be managed in isolation.

These requirements correspond to five of the six operational domains examined in the next section. The sixth, workload management, does not appear as a separate requirement because it cuts across all five. The patent record treats computational load as a control variable rather than as an independent subsystem. Each domain may also support a different degree of autonomy, so a higher-level decision layer is needed to coordinate them and allocate their shared resources.

This coordination layer must anticipate change rather than react only after it occurs. Sea state, weather windows and satellite passes can all be forecast before the relevant decisions must be made. The platform can therefore plan its actions around expected conditions. Autonomy does not emerge by automating each subsystem separately; it depends on coordinating them across the platform as a whole.

4. The Joint Decision Problem of Autonomous Oceanic Computing

An autonomous offshore data center must continuously make decisions, without an operator on site, across six domains (Figure 4). The central argument of this survey is that these domains cannot be managed independently; their close coupling is what distinguishes oceanic computing from earlier offshore facilities.

4.1. Computational Load as a Control Variable

In terrestrial data centers, the relationship between computing and energy is mediated by the electricity grid. Workload schedulers may move or delay jobs in response to the availability and carbon intensity of electricity across geographically distributed sites [18, 94]. A data center may also operate as a flexible load, increasing or reducing its demand to help balance the wider power system [99]. In both cases, however, the facility remains a consumer downstream of the grid. It can change how much electricity it uses in response to external signals, but it does not directly control a generator. Generator output is determined by system-level dispatch, which balances supply and demand across many users.

An offshore platform that harvests its own energy has no grid between the generator and the computing equipment. A patent filed in 2019 describes a system that adjusts the rate of computational work according to the energy available from a variable source. More importantly, it regulates the resistance that a turbine presents to the moving fluid by changing the electrical power consumed by the computer network [17]. This second feature is crucial. Changing the computing load changes the electrical load seen by the generator and therefore shifts the operating point of the energy-conversion system. Computation is no longer simply a consumer of electricity; it becomes a means of controlling the power plant itself.

Two features of the marine environment make this coupling practically valuable. The first is the highly variable nature of wave power. Wave energy arrives as a narrow-band stochastic process, with output oscillating over each wave cycle and changing more slowly with the sea state [58]. The difference between peak and average power can therefore be extreme. At a given sea state, instantaneous power has been reported to reach about 20 times the average [100]. If the power-take-off system generates electricity only during the upward stroke, the peak can reach 60 times the average [101, 102]. The conversion chain must then be able to handle 60 units of power even though it delivers only one unit on average. This low utilization is identified as a major obstacle to the economic viability of wave-energy conversion [102].

The main benefit of load flexibility is not necessarily greater energy capture. The largest peaks contribute relatively little to the total energy available over time. One curtailment study found that a converter rated 87% below the maximum peak could still capture 91% of the available energy. In other words, clipping the highest peaks sacrificed less than 9% of the annual yield [100].

The larger opportunity lies in reducing the size of the machinery. If the computing load can rise when instantaneous generation is high and fall when it is low, the generator and power-electronic chain need not be designed around rare peaks. The same wave-energy literature identifies the avoidance of generator overdesign as an important source of

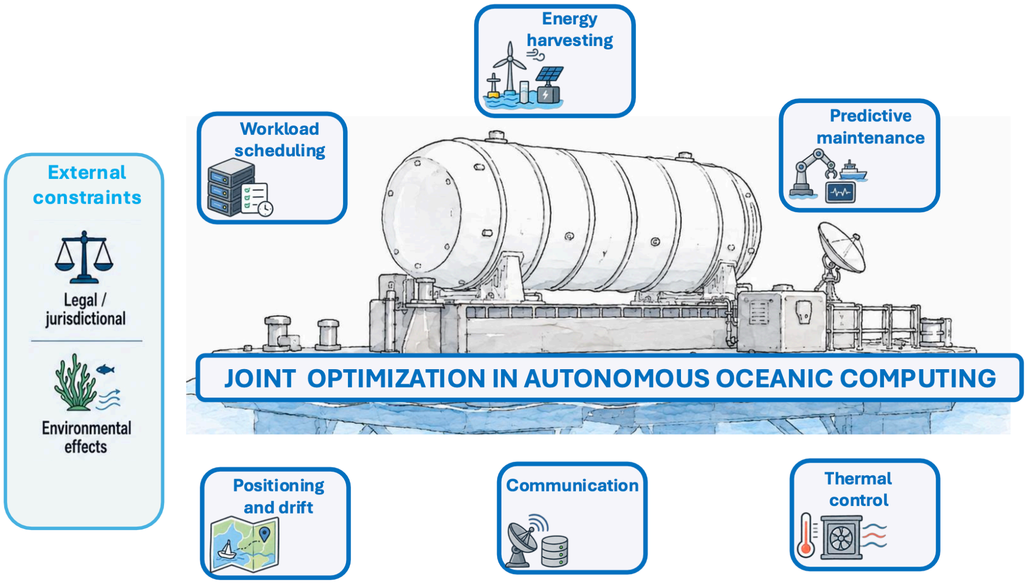

Figure 4: Joint allocation problem of an autonomous oceanic computing platform. Six coupled decisions share the same energy and time resources, namely energy harvesting [17], workload scheduling [18, 94], thermal control [8], communication [21], predictive maintenance [15], and positioning [20, 24]. The platform is also constrained by the legal rules associated with its location [49] and by environmental effects [22].

cost reduction [100]. On an offshore platform, this advantage is especially valuable because mass, capital cost and maintenance access are all tightly constrained. Capturing nearly the same amount of energy with a smaller conversion chain may be more useful than capturing a small additional fraction with larger and heavier equipment.

The second feature is the lack of an alternative destination for surplus electricity. A grid-connected data center that reduces its computational load can export unused power to the grid. An off-grid platform cannot. It must either consume the electricity immediately, store it or curtail its production. Storage is not cost-free: batteries add mass and capital cost, have a limited cycle life and may require vessel-based replacement [24]. Flexible computation therefore gives the platform another way to absorb short-term variations without relying entirely on batteries.

This mechanism complements the smoothing obtained by distributing wave-energy converters across an array [24, 56]. Spatially separated converters encounter waves at different phases, so their combined output is less variable than that of a single unit. Array design therefore smooths the supply through spatial averaging. Load modulation acts on the other side of the balance by reshaping demand over time. The two mechanisms are independent but potentially complementary: one makes generation smoother, whereas the other makes computing demand more responsive.

These approaches have not yet been designed together. Studies of wave-energy arrays optimize the number and arrangement of converters while treating the consumer as a fixed load or a grid connection [56, 57]. By contrast, workload-scheduling studies adapt computing demand while treating renewable generation as an external process that the scheduler cannot influence [18, 94]. A joint design would choose array geometry, load flexibility and storage capacity together, subject to a target for completed computational work. Determining how much physical smoothing can be replaced by flexible demand—and how much storage is still required—remains an open problem.

This coupling raises three further questions. The first concerns the control objective. Terrestrial workload schedulers generally minimize electricity cost or carbon emissions while meeting service-level requirements. An off-grid platform has no electricity tariff to guide its decisions. Carbon pricing applies to maritime transport through the European Union Emissions Trading System, and a global mechanism has also been proposed under the IMO Net-Zero Framework [103, 104]. These mechanisms target greenhouse-gas emissions from ships. For a platform powered entirely by harvested renewable energy, they do not currently provide a direct criterion for deciding which computations should run.

The objective must therefore come from within the platform. One possibility is to maximize the amount of computational work completed per unit of harvested energy. This formulation, however, requires a common measure of "work" across heterogeneous jobs. A training task, an inference request and a data-processing pipeline may consume similar energy while producing very different forms of value. The chosen measure will therefore determine which workloads the scheduler favours. Defining this measure is not simply an accounting issue; it shapes the operating priorities of the platform.

The second question concerns stability. The computing system influences the converter, while the converter's output limits how much computing can be performed. Together,

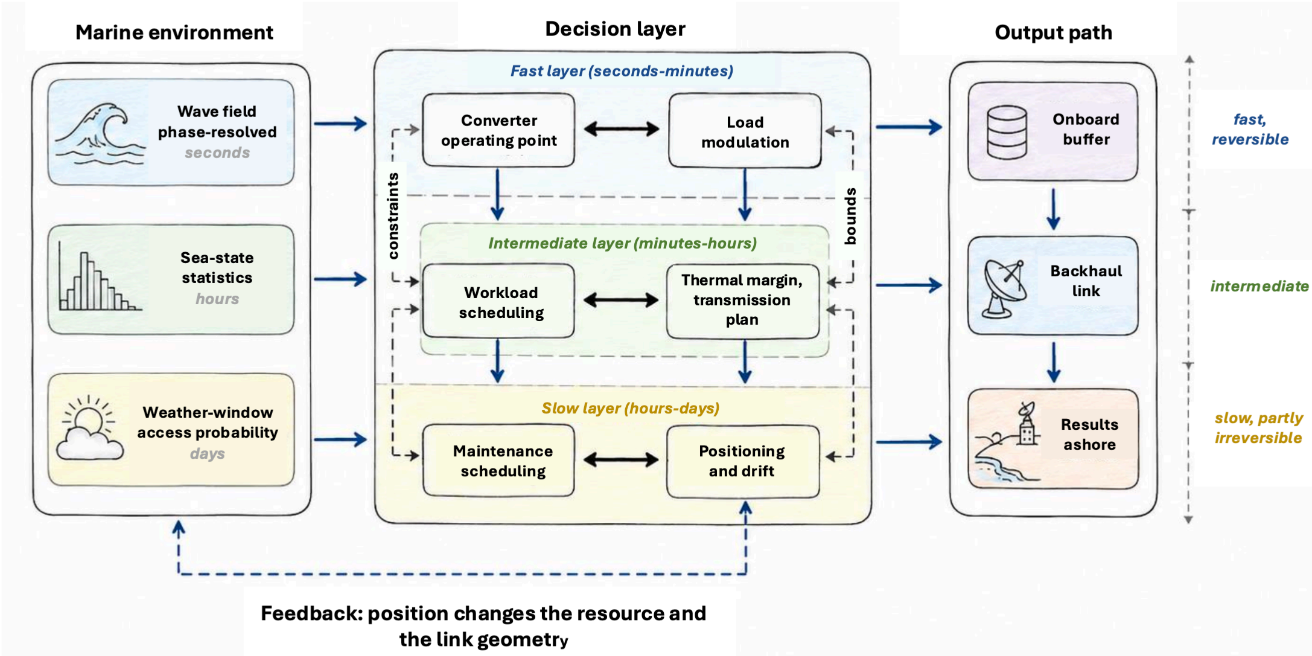

Figure 5: The joint allocation problem of an autonomous oceanic computing platform. Forecasts at three horizons feed three decision layers, which are coupled vertically, since a decision at one scale constrains the layers below and is bounded by those above. The horizons are the wave-by-wave prediction [95, 96], the sea-state forecast [97, 98], and the weather-window statistics that govern physical access [63]. The horizontal coupling between converter operating point and load modulation is the one the patent record claims and the literature has not characterized [17, 58]. The feedback indicates that a change of position alters the resource, the link geometry and the applicable legal regime [20, 49].

they form a closed control loop operating over at least two timescales. Processor frequency and power consumption can be adjusted within milliseconds, allowing rapid load modulation that can follow individual wave dynamics [105]. Workload admission is slower. Once a task begins, it reserves memory, moves data and creates state that cannot be released or reversed without overhead.

The wave-energy converter may also be governed by its own predictive controller [95, 96]. It is not yet clear whether the slower workload scheduler can treat rapid load modulation as a small external disturbance, or whether the computing and converter controllers must be designed as one system. Poorly coordinated controllers could respond to each other in ways that cause oscillation, reduce energy capture or violate workload requirements. The effect of rapid changes in computational demand on the converter's predictive control has not been characterized.

The third question concerns the output side of the platform. Communication capacity varies independently of energy availability. A period of strong waves and abundant electricity may occur when the satellite link is weak or unavailable. Conversely, a good transmission window may arrive when little energy is being harvested. The platform must therefore coordinate energy storage, computation and data buffering across two separate sources of uncertainty.

Storage capacity should consequently not be determined only by the immediate working set of the computing tasks. It must also reflect the joint statistics of energy supply and link availability: how long low-generation periods last, how often communication windows occur and how much computation can be completed before results can be sent ashore. The batteries and data buffers are therefore part of the same scheduling problem as the generator and the workload.

4.2. Degrees of Autonomy and the Joint Allocation Problem

Autonomy is not an all-or-nothing property. Maritime research already provides a scale for describing how much decision-making remains with a human operator and how much is delegated to the system. The scale originated in analytics maturity models [106], in which each level answers a different question, and was later extended in the maritime literature by adding an autonomous level at the top [15]. At the descriptive level, a system reports its current state and answers, "What is happening?" At the diagnostic level, it identifies the cause and answers, "Why is it happening?" At the predictive level, it estimates what will happen next. At the prescriptive level, it recommends what should be done. At the autonomous level, it acts on that recommendation without waiting for a decision from shore. The first four levels therefore support a human decision-maker; only the fifth closes the loop between understanding and action.

A similar progression appears in structural engineering, where ships and offshore structures are evolving into self-aware cyber–physical systems. Continuous monitoring replaces periodic manual inspection, while prediction and planning replace intervention only after a fault has occurred [16]. We adopt this scale because it separates knowing from acting, a distinction that is especially important for an unattended facility at sea. Moving from prescription to autonomy means that the platform must change a setting, operate equipment or move through the physical environment without first receiving approval from shore.

Applying a single autonomy label to an entire facility would hide important differences among its subsystems. We therefore assign a level separately to each of the six decision domains. A submerged pod, for example, may be autonomous in thermal control. Sensors monitor the temperature of the internal cooling loop, and pumps adjust the flow to keep it within the required range. Because the surrounding seawater provides the heat sink, this control can be performed locally. Maintenance may nevertheless remain only descriptive: sensors can detect and report degradation, but any repair still requires external support. A grid-connected facility may reach the prescriptive level in workload management while having no positioning domain at all.

Figure 6 applies this scale to documented systems using only the capabilities reported in their sources, rather than those that the underlying technology might support in principle. Reading across the rows shows that no system reaches a high level of autonomy in every domain. Maintenance and positioning also differ fundamentally from the other domains. A workload, thermal or communication decision can often be implemented by changing a software or control setpoint. A maintenance or positioning decision requires moving mass through water—whether a robot, a vessel, a replacement module or the platform itself. This consumes energy and time that cannot be recovered once the action has begun.

Three patterns emerge from the mapping. First, the highest autonomy levels occur where action remains largely digital. Workload management already reaches prescriptive and autonomous levels in terrestrial data centers because the decision is made and executed in software [18, 94]. Maintenance and positioning generally remain at lower levels because acting on them requires physical movement in a difficult environment.

Second, progress has occurred independently within separate research communities. Data-center engineering develops autonomous thermal control, maritime control develops autonomous positioning, and field robotics develops autonomous inspection and maintenance. A recent surveillance architecture illustrates this separation. It combines a fixed hydrophone with a second hydrophone carried by a mobile robot, whose path is planned to maximize the observability of an acoustic source [26]. This is a prescriptive capability in the maintenance domain, but it was developed without reference to the computing payload that the robot is intended to protect.

Third, no documented system is reported to coordinate decisions across all six domains. This is the main gap addressed by the survey. The energy–compute coupling described in Section 4.1 shows why domain-by-domain autonomy is not enough. The energy available to the platform determines which workloads can run, but the computational load can also change the operating point of the energy converter. Adding autonomous subsystems therefore does not automatically produce an autonomous platform. The constraints that ultimately bind the facility are those that cross domain boundaries.

These interactions can be expressed as one joint allocation problem. At any moment, the platform has a certain amount of harvested power and stored energy. It must divide these resources among computation, cooling, data transmission, auxiliary systems and, where applicable, thrust. At the same time, it must decide which computing tasks to execute, defer, checkpoint or discard; when to transmit stored results; whether to change position; and when to request physical intervention. These choices cannot be separated. Each draws on the same limited energy and resource budget, and each changes the conditions under which the others must operate.

The problem becomes more manageable because its variables evolve on widely separated timescales, as summarized in Figure 7. The fastest dynamics follow the wave period, on the order of seconds. At this scale, the computing load and converter operating point act directly on one another [17, 58]. Communication capacity varies over seconds to minutes, including periodic degradation during satellite handovers [71]. Sea states persist for hours and determine the available wave resource, the thermal margin and the workload mix that can be supported. Weather windows operate over hours to days and determine whether a vessel can safely approach [63]. Changes in the wider resource field and season influence repositioning over days to weeks, while degradation and hardware obsolescence drive maintenance and reconfiguration over months to years. The full problem therefore spans several orders of magnitude, from individual waves to the lifetime of the computing hardware.

Each timescale is usually studied by a different research community, and no single field covers the platform as a whole. A decision made at one scale constrains decisions at shorter timescales and is itself bounded by decisions and conditions that evolve more slowly. The platform must therefore coordinate rapid control actions with longer-term operational commitments.

Figure 5 makes this structure explicit. Forecasts at three horizons enter from the marine environment and feed decision layers operating on different timescales. The layers are coupled vertically: a slow decision sets limits for faster decisions, while fast operation must remain consistent with longer-term commitments. Two features distinguish this structure from a conventional one-way control hierarchy.

The first is the horizontal coupling at the fastest layer. The converter operating point affects the electrical power available to the computing load, while changes in the load alter the operating point seen by the converter. This bidirectional relationship is claimed in the patent records [17], but no published study has yet characterized its dynamics or stability. The second is a return path from the slowest decision to the fastest input. Repositioning changes the marine-energy resource available to the platform, the geometry and quality of its communication link, and the legal regime under which it operates. Position is therefore not merely an output of the control system. By changing the environment encountered by the platform, it also changes the inputs to the fast control layer. The overall structure is consequently a loop rather than a simple cascade: the slowest variable re-enters the system at the point where the fastest one is determined.

This separation of timescales suggests a control hierarchy that the existing literature has not yet assembled. At the fastest scale, the problem is one of power-electronic control, linking load current to the operating point of the wave-energy converter. At the intermediate scale, it is a scheduling problem under two uncertain processes: stochastic energy supply and stochastic capacity for sending results ashore. Carbon-aware and renewable-aware scheduling provide a starting point [18, 94], but their objectives must be adapted. For an off-grid platform, a more relevant objective is the computational work completed per unit of harvested energy, subject to service deadlines. At the slowest scale, the problem becomes one of asset management under restricted physical access, where maritime digital-twin methods are most applicable [15, 16].

Two structural features distinguish this problem from its terrestrial counterpart. The first is the absence of an external reference signal. A grid-connected facility can optimize against electricity prices or carbon intensities set by the wider power system. An off-grid platform has only its own state, commitments and forecasts. Its objective must therefore be defined internally, and the forecast is not simply information supplied to the controller; it is part of the control loop itself. The second feature is the asymmetry between reversible and irreversible decisions. A scheduling choice can often be changed in the next interval. A repositioning decision consumes propulsion energy that cannot be recovered, and a maintenance decision commits a vessel, personnel and a weather window. The platform therefore combines a fast, frequently revised control layer with a slow layer whose actions have lasting consequences. Mathematically, this resembles a stochastic control problem with switching costs, in which changing the operating state incurs a non-recoverable cost. The oceanic-computing problem has not yet been formulated in these terms.

4.3. Maintenance and Positioning

4.3.1. Prediction without Access

Maritime engineering has made substantial progress in diagnosis and prediction. Continuous monitoring of structures, machinery and environmental loads is replacing periodic inspection with condition-based and predictive maintenance [15, 16]. In principle, these capabilities can progress from diagnosis and prediction to recommended actions and, eventually, autonomous intervention. Offshore, however, knowing what should be done is not the same as being able to do it.

A model may predict that a component will fail within a given period, yet the platform may be unable to intervene. Physical repair requires a vessel, and vessel access depends on suitable weather that follows its own schedule. This changes the purpose of prediction. At a terrestrial facility, prediction mainly provides time to repair a component before it fails. At the offshore site examined in [63], conventional crew-transfer vessels were unavailable for roughly one-third of the year, and winter conditions sometimes produced several consecutive days without a safe operating window. Prediction therefore serves mainly to organize maintenance around the limited opportunities for access.

The platform must determine which interventions can be combined during the next suitable weather window and which forms of degradation can safely be tolerated until then. The useful predictive horizon is therefore determined not only by sensor and model performance, but also by the access statistics of the site. A highly accurate prediction has limited operational value if it extends only a few hours ahead but the expected wait for a vessel-access window is several days.

Existing systems address this constraint in two ways. The first is to reduce or eliminate the need for access. A sealed submerged pod is designed to operate without human intervention throughout its deployment, but this reliability comes at the cost of reconfiguration: failed or obsolete hardware cannot be replaced until the entire pod is recovered [14, 34]. The second approach is to keep inspection capability permanently on site. This arrangement has been demonstrated for acoustic surveillance, in which a mobile robot carries a sensor around the facility and plans its path to maximize the observability of the event being monitored [26]. A resident robot turns part of maintenance from a vessel-access problem into an onboard control problem. Among the systems reviewed in Figure 6, this is the only reported route by which maintenance advances beyond diagnosis without first requiring a vessel. To our knowledge, no existing design combines the long unattended operation of a sealed pod with a resident robotic system capable of inspection and intervention.

4.3.2. Position as a Decision

For a moored or anchored facility, position is fixed. For a self-propelled platform, it becomes an operational decision whose energy cost was established in Section 3. In deep international waters, where the seabed may be beyond reach and anchoring may not be available, holding station is the active and energy-intensive state. Drifting is the passive and lower-energy state. The main question is therefore not simply where the platform should be, but whether a particular location is valuable enough to justify the propulsion energy required to reach and maintain it.

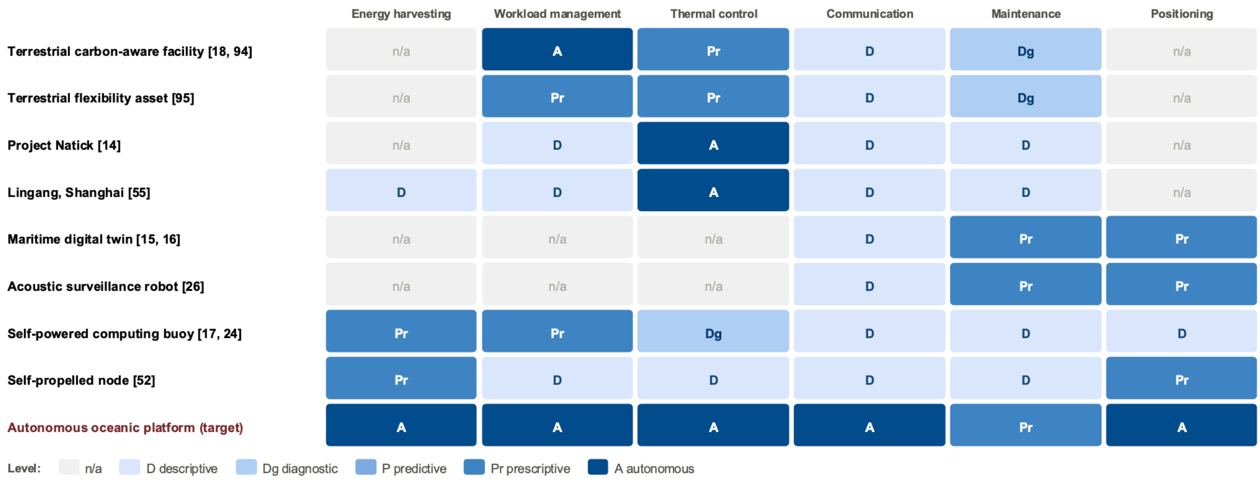

Figure 6: Reported level of autonomy of documented systems across the six decision domains, on the capability scale of [15]: D descriptive, Dg diagnostic, P predictive, Pr prescriptive, A autonomous. "n/a" indicates a domain that does not exist for the configuration. Levels reflect the capabilities reported in the cited sources, rather than those that the technology could support in principle.

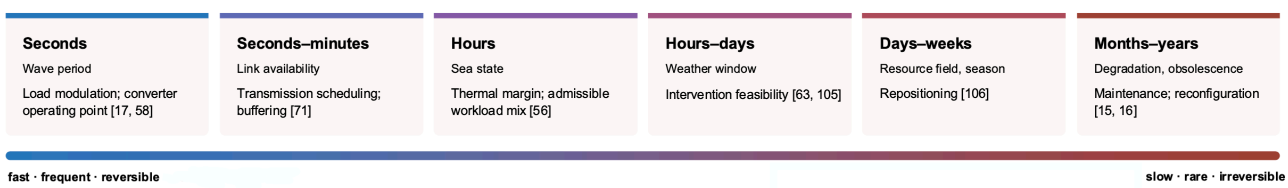

Figure 7: Timescales of the decision variables on an autonomous offshore computing platform, the domain in which each is set, and the sources characterizing it.

The evidence reviewed here identifies four reasons why relocation may be worthwhile. First, the wave-energy resource varies greatly with latitude. Reconstructions based on several decades of historical wave conditions show values ranging from approximately 2 kW per metre of wave front near the equator to 80 kW per metre in the Southern Ocean [107]. Moving can therefore give the platform access to a substantially richer energy resource.

Second, marine heat waves reduce the temperature difference that drives heat from a submerged facility into the surrounding water. When this cooling margin becomes too small, migration to cooler waters has been proposed as a possible response [20]. Third, communication quality depends on the platform's location, orientation, motion and geometry relative to the communication terminal or satellite, and this dependence has been quantified for optical links [82]. Fourth, the probability of gaining physical access for inspection or repair also varies by site because weather and sea conditions determine when a vessel can approach [63].

A mobile platform can respond to all four factors, but they may favour different locations. A region with stronger waves may provide more energy but poorer communication or more difficult maintenance access. Cooler water may improve heat rejection without providing the best satellite geometry. The relative value of each factor also depends on the workload. A task with a strict delivery deadline may make connectivity the dominant consideration. A task that can be delayed may instead favour a location with abundant renewable energy, even if communication opportunities are less frequent.

Position therefore cannot be optimized independently of computing. The value of a location depends on the amount and type of computational work it enables. Moreover, the energy spent reaching and holding that location is no longer available for the workload itself. Position consequently belongs to the same joint allocation problem described in Section 4.2. Position is also unique among the decision variables. Changing it may take days, and the propulsion cost must be paid before the benefits of the new location are fully known.

4.4. Forecasting as Part of the Control Loop

Section 4.2 showed that an off-grid platform cannot rely on an external signal, such as an electricity price, to guide its operation. Instead, it must make decisions using forecasts of its own environment. These environmental forecasts span three distinct time horizons, each using different methods and supporting a different layer of the platform's control system.

At the shortest horizon, the sea surface is predicted wave by wave, with the phase of each wave explicitly resolved [95, 108]. Model predictive control of a wave-energy converter typically requires a forecast of the incoming wave-excitation force about 30 seconds ahead [95]. One approach treats surface elevation as a time series and predicts its future values from its recent history. Using this method, a linear autoregressive model has been reported to forecast ocean swell accurately for approximately two wave periods [97]. The prediction requirements for real-time converter control have also been quantified separately [96].

Obtaining suitable measurements is the main practical challenge. X-band navigation radar has been validated in open water as a sensor for producing these short-term forecasts. The potential gain is substantial: predictive control informed by wave forecasts has been reported to increase captured power by more than 50% compared with causal controllers, which respond only to conditions that have already been observed [95].

This shortest horizon is also the one at which computational load can act on the energy converter. The energy–compute coupling described in Section 4.1 operates over tens of seconds and therefore inherits both the capabilities and the limitations of deterministic wave prediction. A workload whose processing rate can be changed within this interval can be coordinated with the approaching wave train. A workload that can be adjusted only at a coarser timescale cannot respond quickly enough.

At the intermediate horizon, the nature of the prediction changes. Operational wave models no longer attempt to predict the phase of every wave. Instead, they describe the statistical state of the sea, providing quantities such as significant wave height and mean wave direction over intervals of one to three hours [98]. Wave-by-wave prediction is neither available nor necessary at this range. These forecasts support slower decisions, including which workloads can be admitted, how much thermal margin is available and when stored results should be transmitted. Such decisions depend on the expected distribution of conditions rather than on the arrival of an individual wave.

At the longest horizon, meteorological and oceanographic, or metocean, forecasts support the planning of physical actions. Weather-window analysis for offshore operations uses hourly reanalysis and forecast data to estimate the probability that conditions will remain suitable for a specified period. The required duration depends on the intervention: personnel transfer, inspection, lifting and major repair each need different access windows [63]. These forecasts guide maintenance and repositioning decisions, and their accuracy determines how far in advance the slow and partly irreversible actions discussed in Section 4.3 can be planned.

The three forecasting layers have largely been developed by separate research communities. Phase-resolved prediction is studied mainly for the control of individual wave-energy converters, with the objective of maximizing captured power. Sea-state forecasting is developed for marine operations, whereas weather-window analysis is used for offshore asset logistics. These methods have not yet been combined into a common forecasting architecture.

An autonomous computing platform would bring all three forecasts into one decision system. Their errors would then propagate into different operational domains and have consequences of very different scales. Intuitively, an error in a 30-second forecast may cost converted energy. An error in a three-hour forecast may lead the platform to admit an unsuitable workload or adopt the wrong transmission plan. An error in a three-day forecast may result in an unnecessary or unsuccessful vessel mobilization.

An important open question is therefore how uncertainty at each forecasting horizon propagates through the joint allocation problem. Answering it would reveal which forecasts most strongly affect system performance and where improvements in sensing, modelling or computation would provide the greatest value. Such an analysis is a prerequisite for any quantitative assessment of autonomous oceanic computing.

4.5. Lessons Learned

Three conclusions are summarized from this section. First, software-based decisions, such as workload scheduling, can already reach prescriptive or autonomous levels. Maintenance and positioning are less advanced because acting on them requires physical intervention or movement through water. Second, progress remains fragmented across data-center engineering, maritime control, field robotics and forecasting. Third, no documented system coordinates all six domains.

These domains share the same limited resources but operate over timescales ranging from seconds to years. Fast decisions, such as load modulation, are frequent and reversible. Slow decisions, such as relocation or hardware replacement, have lasting effects and reshape the conditions under which faster decisions are made. Moving the platform, for example, changes its energy supply, communication link, cooling conditions and applicable operating regime. Unlike a grid-connected facility, an off-grid platform also lacks an external signal such as an electricity price, so its objective must be derived from its own state, workload commitments and forecasts.

Autonomy within individual subsystems therefore does not amount to autonomy of the complete platform. Oceanic computing requires a higher-level coordination layer that allocates shared resources across all six domains and plans ahead using forecasts of sea state, weather windows and satellite passes.

5. The Legal and Regulatory Regime

The legal framework governing an oceanic computing platform changes with its location. The siting axis in Figure 2 is therefore also a jurisdictional axis. Two distinct legal regimes apply: one governs the platform as a physical object, while the other governs the data it processes. Each imposes its own environmental and compliance obligations, but the two regimes arise from different bodies of law and do not align neatly. The interaction among these regimes is least

Table 3
Legal rules applicable to each siting option of Figure 2.

Siting	Applicable law, platform	Applicable law, data	Unsettled
Coastal / nearshore	Full coastal-State authority, domestic permitting	Domestic data law of the coastal State	Little. Ordinary offshore practice applies
Exclusive economic zone	UNCLOS Arts. 60, 80. Authorization, safety zone, removal duty	Law of establishment and of the data subject	Whether a mobile node is an installation
High seas	UNCLOS Art. 87 freedom to construct, Art. 60(8) no island status	Unchanged. Obligations follow operator and data subject	Ship or installation. Removal standard. Safety-zone basis

Table 4
Maritime zones under UNCLOS and their implications for a computing platform.

Zone	**Coastal-State authority**	**Articles**	**Implication for the platform**
Internal waters, territorial sea	Full sovereignty, subject to innocent passage	2–32	Ordinary domestic regulation applies. Permitting, safety and environmental law of the coastal State govern the facility
Contiguous zone	Enforcement limited to customs, fiscal, immigration and sanitary matters	33	No general regulatory power over the platform. Control is confined to the four listed matters
Exclusive economic zone, continental shelf	Exclusive right to construct, authorize and regulate artificial islands, installations and structures, and jurisdiction over them	60, 80	Coastal-State authorization required. Safety zones of up to 500 m may be established. Under Art. 60(8) the structure has no island status, no territorial sea of its own, and no effect on maritime delimitation
High seas	None. No State may claim sovereignty	87–89	Freedom to construct installations, subject to due regard for the rights of other States. Reserved for peaceful purposes. The Convention's detailed rules on installations apply only by analogy

settled for emerging far-offshore configurations. Table 3 summarizes the rules applicable to each siting option.

5.1. The Platform under the Law of the Sea

The United Nations Convention on the Law of the Sea (UNCLOS) divides the ocean into zones in which coastal-State authority progressively declines: internal waters, territorial sea, contiguous zone, exclusive economic zone and high seas [49]. Table 4 summarises each zone, the powers available to the coastal State and the consequences for a computing platform. The submerged deployments discussed in Section 3 fall within these established regimes. They are therefore typically developed with national operators and permitted through the same procedures used for conventional offshore infrastructure. UNCLOS also makes clear what such a structure does not become: regardless of its scale or permanence, a computing platform creates no jurisdictional zone of its own. Beyond national jurisdiction, however, the framework is less certain. The high-seas regime is permissive in principle and generally allows such structures, provided that other lawful uses of the sea are respected. Yet it offers little guidance tailored to computing platforms, leaving their legal treatment only partly defined [109].

The important uncertainty is classification. A self-propelled node that generates its own power, moves under its own control and carries neither cargo nor passengers fits poorly within the categories recognised by UNCLOS. If classified as a ship, Articles 91, 92 and 94 apply: it must be registered, fly the flag of one State and remain under that State's jurisdiction and control in administrative, technical and social matters. If classified as an installation, its construction on the high seas is governed by the freedom recognised in Article 87. However, UNCLOS's detailed provisions for installations were written for the exclusive economic zone and continental shelf and apply beyond those zones only by analogy [109]. This distinction has practical consequences. It determines the applicable safety standards, the authority responsible for inspection, the rules governing collisions with vessels and whether enforcement mechanisms associated with flagged ships are available. A platform operating without a flag may be treated as stateless. Under Article 110, a warship may then board it on the high seas because UNCLOS permits the boarding of a ship without nationality.

Mobility deepens this uncertainty. A fixed installation remains in one legal location throughout its operating life. A self-propelled node can cross the boundaries shown in Table 4, and the applicable rules may change each time it does so. Section 4.3 treats repositioning as an operational choice shaped by energy availability, thermal conditions and connectivity. Legally, however, the same movement can also shift the regime governing the platform. No existing formulation of the allocation problem accounts for this change in applicable law.

Table 5
Data-related regimes and their effect on siting, as of mid-2026.

Jurisdiction	Instrument	Transfer rule	Constrains position
European Union	GDPR [50]	Permitted under safeguards	No
United States	CLOUD Act [51]	Disclosure regardless of storage	No
China	PIPL [110], CSL [111], DSL [112]	Assessment or certification	Yes
Vietnam	Law 116/2025/QH15 [113]	Localization for listed data	Yes
Singapore	PDPA [114]	Permitted under safeguards	No
Japan	APPI [115]	Permitted under safeguards	No
Hong Kong	PDPO, s. 33 [116]	Restriction not in force	No
India	DPDP Act, s. 16 [117]	Permitted unless notified	No

5.2. Jurisdiction over the Data

The second legal regime follows the data rather than the platform, and it operates according to a different logic. Data residency describes where data are physically stored. Data localization requires them to remain within specified borders. Data sovereignty identifies the legal framework that governs them, which need not match either their physical location or a localization requirement [118]. Moving servers beyond national territory therefore does not move the data beyond national law. The General Data Protection Regulation (GDPR) applies to the personal data of individuals in the European Union wherever those data are processed. Moreover, the assessment of foreign-surveillance risks required for international transfers is a continuing obligation [119]. Under the United States CLOUD Act, data controlled by a United States provider is subject to disclosure obligations regardless of where it is stored. An operator remains bound both by the law of its place of establishment and by the law governing its customers' data. Moving the hardware into international waters removes neither obligation.

National approaches to data transfer and localization vary widely, as summarized in Table 5 [120]. Many jurisdictions permit international transfers when appropriate safeguards are in place, and these safeguards can continue to apply as a workload moves. Localization rules are more restrictive because they tie data storage or processing to a defined territory. For a mobile platform, they directly limit where a regulated workload may operate. Singapore and Japan generally allow transfers subject to appropriate safeguards [114, 115]. China and Vietnam impose localization requirements on particular categories of data or operators [110, 113]. Hong Kong and India generally permit transfers by default. Hong Kong's statutory transfer restriction has never been implemented [116], whereas India allows transfers except to countries formally notified as restricted [117].

Offshore deployment also creates distinct security concerns because a platform may operate without continuous physical supervision. This makes attacks requiring direct access, such as cold-boot recovery of data from memory, more relevant [121]. Stronger protection of cryptographic keys and remote attestation are therefore needed to verify platform integrity without an on-site inspection [122]. A secure design keeps key-management functions onshore so that gaining physical access to the platform does not reveal readable data. The same separation does not remove privacy obligations. Under the GDPR, operators must implement data protection by design, minimize the personal data they process and apply security measures appropriate to the risk [123]. Cross-border transfer rules further constrain where personal data may be stored or processed.

Two consequences follow. First, an offshore computing platform is not a data haven. Treating it as one misreads the applicable law because legal obligations follow the operator and the data subject, not simply the location of the machines. Second, mobility makes position legally consequential. A workload subject to localization requirements restricts the places in which the platform can operate. The allocation problem in Section 4.2 must therefore include legal constraints alongside physical ones. A deferrable job with no jurisdictional restriction may follow the most favourable energy resource, whereas a regulated workload cannot leave the zone in which it is permitted to be processed. In this sense, position becomes a compliance variable.

5.3. Environmental Obligations

Part XII of the UNCLOS establishes a general duty to protect and preserve the marine environment [49]. Article 208 requires States to adopt laws that control pollution from installations under their jurisdiction, while Article 206 requires an assessment when an activity may cause substantial pollution or significant harmful changes. Within national jurisdiction, these duties are implemented through domestic environmental-assessment procedures. The Baltic Sea illustrates the practical complexity that can arise when several national regimes overlap within a semi-enclosed sea [54].

Three effects specific to computing platforms fall within these provisions and are examined in Section 6: the discharge of heat into surrounding water, underwater acoustic emissions and the fate of the structure at the end of its life. Article 60(3) requires installations that are no longer in use to be removed, with due regard for navigational safety and the marine environment. No equivalent specific removal duty applies on the high seas, where only the broader obligations under Part XII remain. For a sealed pod whose operating life may be measured in years while its structure may last for decades, decommissioning is therefore a significant obligation. Yet neither the disclosures examined in Section 3 nor the other documents reviewed explain how recovery or removal would be managed.

5.4. What Remains Unsettled

Four questions remain unresolved. First, should a self-propelled computing platform be classified as a ship, an installation or a third legal category, and what would each classification mean for registration, inspection and liability? Second, is there any high-seas equivalent to the safety zones permitted in the exclusive economic zone, and can such a platform therefore keep approaching vessels away? Third, how should a fleet of platforms that occupies an area and moves within it satisfy its duty of due regard under Article 87? Fourth, how should the maritime legal regime interact with data-protection obligations when the two impose conflicting limits on where processing may take place?

These questions will need to be addressed before large-scale commercial deployment. Their answers will shape which oceanic-computing configurations are commercially viable as much as the technical constraints examined in this paper.

6. Environmental Effects

Section 5 examined which State has authority over an oceanic-computing platform and the data it processes. Here, we turn to the platform's effects on the marine environment and the standards used to assess them. UNCLOS links these issues because it governs both jurisdiction and environmental protection. The environmental case for offshore computing rests partly on the impacts it avoids: land occupation, freshwater consumption and reliance on grid electricity generated elsewhere [1, 5]. This section examines four effects created by placing computing infrastructure at sea. Evidence exists for each, but comes largely from studies of other industries. The final subsection identifies what has yet to be measured.

6.1. Thermal Discharge

Data centers generate substantial heat, and an offshore facility ultimately releases it into the surrounding water. The closest established analogue is cooling at coastal thermal and nuclear power plants, where a large share of primary energy is discharged to the sea as waste heat [124]. Continuous discharge from such plants has been reported to raise local water temperatures by 4-7 °C, creating surface plumes that spread away from the site [125]. Reported effects include lower dissolved oxygen, changes in benthic community structure and biodiversity, suppressed phytoplankton growth within the plume, thermal stress in sensitive organisms and altered stratification of the receiving water [124–126].

The analogy applies to mechanism, not scale. A nuclear plant rejects heat at the gigawatt scale, whereas the largest disclosed offshore computing facility is rated in tens of megawatts [55]. Even so, the underlying processes and the mitigation strategies developed for them remain relevant. Discharge depth is the principal design variable. Releasing heat below the mixed layer can improve dilution and reduce effects on photosynthetic communities [126]. In the structural regimes shown in Figure 2, however, this choice is constrained by platform design. A submerged pod releases heat at depth, whereas a floating platform releases it into the upper water column. Because this layer is illuminated and biologically active, surface discharge may expose planktonic communities more directly to thermal perturbation. Deeper discharge may reduce this exposure, although the ecological outcome depends strongly on the release depth and local oceanographic conditions.

Mobility adds a question with no direct counterpart in the power-plant literature. A mobile platform spreads its thermal discharge along its route rather than concentrating it at a fixed site. This can reduce local exposure but enlarge the total affected area. Whether that represents an environmental improvement depends on whether ecological damage occurs only above a threshold or accumulates over time. This question has not yet been examined for a moving marine heat source.

6.2. Underwater Noise

Sound from an underwater source can remain detectable over much greater distances than sound from a comparable source on land [127]. Human-generated underwater noise is therefore recognized as a stressor on marine fauna. Its main effect is acoustic masking: background noise reduces the distance over which animals can communicate, navigate, forage and detect predators [127, 128]. The International Maritime Organization has issued revised guidelines for reducing underwater noise from shipping [129], and its Marine Environment Protection Committee has approved an accompanying action plan [130]. Offshore computing platforms, however, have a different acoustic profile and do not fit neatly within this framework.

Noise from ships is dominated by propeller cavitation and flow around the hull during transit [131]. A moored or slowly moving computing platform produces little of either. Its main source is instead likely to be structure-borne noise from pumps, compressors and cooling equipment. Vibrations from this machinery enter the hull and radiate into the surrounding water through a pathway that has been characterized experimentally for ship machinery [131, 132]. This noise differs from conventional shipping noise in two respects. First, it is continuous rather than episodic because the machinery operates whenever the facility is computing. Second, it is concentrated at discrete frequencies determined by the rotational speeds of the equipment and its blades.

This mechanism also links environmental noise to physical security. External narrowband sound can pass through the enclosure of a submerged pod and disrupt the storage devices inside [26, 48]. The same structural pathway operates in reverse, allowing vibrations from internal machinery to pass through the hull and radiate into the water. The hull is therefore a two-way acoustic interface. Measures that reduce vulnerability to external acoustic injection may also reduce noise emitted by the platform, but no study has yet examined both directions together.

6.3. Colonization and the Artificial Reef Effect

Submerged structures are colonized by marine organisms. Section 3 treated the resulting biofouling as a thermal liability, but the same process also has ecological consequences. The patent record presents colonization as a design objective, describing the vessel as an artificial reef intended to support the surrounding ecosystem [34]. Colonization follows a well-documented succession and depends on the construction material, surface properties and local environmental conditions [22, 69]. Yet offshore practice has no established criterion for balancing thermal performance against ecological effects, leaving this trade-off largely unexplored. Antifouling measures preserve heat-transfer performance by suppressing colonization, whereas allowing organisms to settle can support local biomass but reduce cooling efficiency. A facility whose thermal margin has already been narrowed by marine heat waves [20] has limited flexibility in making this trade-off. In practice, the outcome is determined implicitly by the coating specification rather than through a deliberate environmental decision.

6.4. End of Life and the Measurement Gap

Removal obligations apply when offshore installations are no longer in use. These duties are codified for the exclusive economic zone and continental shelf, with a similar expectation elsewhere [49]. The marine structures considered here are designed to remain in service for decades, whereas their computing hardware may be replaced within a few years. Decommissioning is therefore a foreseeable event within any operator's planning horizon. Yet none of the disclosures reviewed in Section 3 explains how a platform would be recovered. Moreover, the programme responsible for the most extensively documented deployment was closed for commercial, rather than environmental, reasons [14].

More broadly, the available evidence is indirect. Each effect discussed above is supported by research from another industry: thermal discharge from power generation, underwater noise from shipping, colonization of marine infrastructure, and decommissioning in the oil and gas sector. None has been measured directly at an offshore computing facility. Environmental assessments conducted for the deployments reviewed here are not publicly available, and the performance figures released by operators concern energy and water use rather than effects on the marine environment [55]. Establishing environmental baselines and monitoring protocols specifically for this class of facility is therefore essential to support the environmental claims on which the paradigm partly rests. At present, such evidence is absent.

The physical installation creates three further impacts. First, submarine power cables produce electric and magnetic fields during operation. These fields may affect species that use magnetoreception or electroreception for orientation and navigation [133, 134]. Benthic organisms may face particularly persistent exposure because cables are commonly buried within, or placed close to, the seabed [134]. Electrical losses can also heat the surrounding sediment, although field measurements from operating installations remain limited [135].

Second, corrosion protection can introduce metals into the environment. Galvanic anodes gradually dissolve during operation, releasing aluminium and zinc into the surrounding water. Reported environmental concentrations generally remain below observed toxicity thresholds [136].

Third, the infrastructure carries embodied carbon: the emissions associated with manufacturing and constructing its components. Servers are replaced more frequently than the supporting offshore structure. As operational electricity becomes less carbon intensive, these embodied emissions account for a growing share of the overall footprint [137]. Mitigation measures may be adapted from offshore wind practice, although their suitability for computing platforms must be assessed case by case. Burying power cables increases the distance between the conductor and the seabed surface, reducing—but not eliminating—local electromagnetic exposure [133]. Burial conditions also affect heat dissipation and cable ampacity, or the current that a cable can carry safely, through the thermal properties of the surrounding sediment [138]. German offshore regulation limits sediment warming to 2 K above ambient at 20 cm below the seabed [139]. Meeting this requirement constrains cable size, burial design and current rating [139]. Impressed-current cathodic protection can avoid the continuous metal release associated with galvanic anodes, but it requires an actively powered subsystem and produces chlorine-derived oxidants through electrolysis at the anode [140].

7. Open Challenges and Research Roadmap

The gaps identified in the preceding sections are collected here and organized by the horizon over which they can be addressed. These issues differ in the evidence, models, and technological maturity required to investigate them. Some can be addressed with existing tools. Others require measurements that are not available at the moment, while a third group still lacks the framework needed to formulate the problem. Economic feasibility cuts across these categories. Specifically, the assessment framework exists, but several key inputs remain unavailable.

7.1. Questions Answerable with Existing Methods

Four problems can be addressed with the tools already available, and each would produce a result that the field currently lacks.

The closed loop between load and converter. The patent record claims that computational load can regulate the resistance offered by a turbine to its fluid flow [17], which makes the load and the converter a single dynamical system. That system has never been characterized. Establishing its stability margins and its bandwidth, and determining when load modulation interferes with the predictive control of the converter itself [95, 96], would yield the first control model of a computing platform that acts on its own power source.

Joint design of array geometry and load flexibility. Array smoothing and load modulation reconcile a stochastic supply with a demand through independent mechanisms [17, 56, 57]. The wave-energy literature optimizes the first while treating the consumer as fixed, and the scheduling literature optimizes the second while treating the supply as given. Choosing them together, against a target on delivered work and a constraint on storage, would establish how much array size can be traded for load flexibility, and at what cost.

Propagation of forecast uncertainty. Three forecasting horizons feed one decision-maker, and their errors enter different domains [63, 97, 98]. Tracing how uncertainty at each horizon propagates into the allocation problem would identify which forecasts are worth improving, which is a prerequisite for deciding whether to invest in sensing or in computation.

Levelized cost of computation. The stochastic assessment framework developed for offshore energy can be transferred to oceanic computing once the denominator is replaced [63]. Applying it to assumed inputs would reveal which parameters most strongly influence the outcome and, therefore, what future instrumented deployments should measure first.

7.2. Questions blocked on measurement

Three problems cannot be resolved from the public record as it stands.

Achieved performance of existing deployments. Capacity factors, availability, intervention frequency, and operating cost are undisclosed for every system reviewed here, and the figures that operators publish concern power usage effectiveness rather than delivered computation [14, 55]. Without them, the economic question remains open in principle and not in detail.

Environmental baselines. Thermal discharge, radiated noise and colonization have each been studied for other marine industries, and none has been measured for a computing facility [22, 126]. The distinctive questions, namely the effect of a *mobile* thermal source and of *continuous tonal* machinery noise, require instrumented deployments instead of analysis.

Acoustic reciprocity of the enclosure. External acoustic emissions can couple through a pressure vessel and disturb the storage devices inside [26, 48], and internal machinery radiates outward through the same path. Whether the same design measures reduce both directions can be settled only by experiment on a real enclosure, and no such measurement exists.

7.3. Questions Requiring a New Formulation

Three problems are open in the stronger sense that the framework for posing them does not yet exist.

Cross-domain coordination. No documented system coordinates decisions across the six domains of Figure 6. The problem mixes a fast reversible layer with a slow irreversible one, which places it in the class of stochastic control problems with switching costs. It has not been posed in these terms for a marine computing platform.

Position as a compound variable. Repositioning is simultaneously an energy decision, a thermal decision, a connectivity decision, a maintenance decision and a change of applicable law [20, 49, 80]. No formulation carries all five, and the legal term has no representation in the formulations reviewed here. A platform that drifts adds a further coupling, since its trajectory follows currents, tides and wind, and is therefore forecastable from the same metocean data that drives the energy and access layers [63, 98]. Predicting where the platform will be predicts which satellites it will see and under what pointing geometry, so drift forecasting and link forecasting are intimately connected problems, and neither the maritime nor the satellite literature treats them as such.

Fleet-level architecture. The patent record discloses power sharing across inter-tethered units [24], and optical links between platforms would permit results to be aggregated before transmission [76, 80]. A fleet that shares both power and bandwidth is a different object from a collection of independent platforms, and its control architecture and its economics remain undescribed.

7.4. The Economic Feasibility

Offshore deployment has a different cost structure from onshore deployment. It avoids land-acquisition costs, which are substantial and rising in constrained markets [1]. However, it can add costs for marine engineering, installation vessels and qualification of equipment for the marine environment [14]. On the operating side, it removes freshwater consumption for cooling [5] and reduces the energy devoted to heat rejection, since seawater provides a heat sink whose temperature is more stable than ambient air [8, 9]. It adds intervention costs governed by vessel availability, and not by labour rates. In offshore wind, where similar vessels and weather constraints apply, maintenance fleet selection can change energy costs by more than EUR 14/MWh, from about 124 to 110, while also reducing cost variability [63]. A major source of this cost is downtime caused by limited weather windows and logistics availability. In adverse conditions, these factors can account for up to 60% of total downtime, while a 20% reduction in winter accessibility increases the cost by EUR 9–12/MWh [63].

The offshore energy sector assesses projects through the levelized cost of energy, defined as the present value of lifetime costs divided by the present value of energy delivered, with the denominator computed net of downtime and accessibility restrictions and not as theoretical output [63]. The same structure applies to an offshore computing platform once the denominator is replaced by the present value of computational work completed, net of the time during which the platform cannot compute or cannot deliver results. The substitution is substantial, because the terms that populate the denominator differ from those of a terrestrial facility. Availability is an outcome of the interaction between harvested energy, link capacity, failure occurrence and weather-window access, and not a fixed correction factor [63]. Work completed depends on the energy actually harvested, which is stochastic [56, 58], on the fraction of results that can be transmitted, which is also stochastic [73], and on the propulsion energy diverted to station keeping, which competes directly with computation. A platform can be energized and idle because its link is unavailable, or connected and idle because its supply has fallen, and a terrestrial availability model represents neither state. The appropriate assessment method is therefore stochastic, as offshore wind has already recognized, and the basis of comparison must be delivered work and not installed capacity.

The inputs for such an assessment are missing. The capital cost by component, intervention frequency and cost, achieved availability, and the value of the work performed are absent from the record for every system examined here. The paradigm is therefore being capitalized in advance of a published economic case. The configurations that followed Natick differ from it precisely in removing the grid connection, the cable and the fixed position. Whether that alters the conclusion is the central open question.

7.5. Roadmap

The dependencies among these questions suggest an order. The control-theoretic work in Section 7.1 can proceed immediately and produce the models required for a cross-domain formulation. Instrumented deployment, whether of a full facility or of a representative structure, unblocks the environmental and acoustic questions and supplies the inputs for economic assessment. The formulation problems depend on both. The legal questions are on an independent track and are the slowest, since they are resolved by state practice and interpretation and not by research. Moreover, they may become binding before the technical questions are settled. The subsystems are individually mature. Wave energy conversion, seawater heat rejection, maritime autonomy, and satellite backhaul all have substantial literatures and demonstrated hardware. Therefore, future progress will depend on the coupling between them, on evidence that the resulting system performs, and on a legal status for the object that performs it.

8. Lessons Learned

This survey has examined oceanic computing as a unified infrastructure class. This emerging paradigm encompasses submerged pods, moored platforms, converted vessels, and self-propelled nodes operating in international waters. These configurations are usually studied separately, and the technologies that support them are often considered independent subsystems. Examining them together reveals six broader lessons.

The paradigm predates its recent visibility, and much of its early development is documented in patents. A complete floating, wave-powered, seawater-cooled computing facility was disclosed in 2007 [13], nearly two decades before the first reviews of the subject appeared [23, 25]. Reading that corpus as a source of design requirements recovers a body of engineering that the peer-reviewed literature has not codified [46].

Technical demonstration and commercial viability are separate questions here, and the former does not imply the latter. For instance, sealed deployment reported a component failure rate of about one-eighth that of a comparable land facility, yet the programme was later discontinued for commercial reasons [14]. The underlying trade-off weighs reliability against the ability to change the configuration after deployment, and not against cost [13, 24].

Computational load is a control variable and can be optimized according to the energy supply and storage. The record claims regulation of a turbine through adjustment of the rate of computational work [17], which inverts the direction assumed by the terrestrial scheduling literature [18, 94, 99]. The benefit lies less in additional harvested energy than in the possibility of sizing the conversion chain. It also gives the platform a second smoothing mechanism, independent of the geometric smoothing obtained by dispersing converters [56, 57].

Autonomy in individual domains does not automatically lead to autonomy at the system level. Autonomy is highest when decisions can be implemented through software, as in workload scheduling or thermal control. It is lowest when execution requires physical intervention in the marine environment, as in maintenance [15, 16]. The binding constraints are those that cross domains, and no documented system coordinates across them. Offshore prediction supports scheduling and not prevention. The useful predictive horizon therefore depends on site accessibility and the availability of intervention windows [63].

Position is expensive and legally consequential. On the high seas, mooring may be unavailable or impractical. A platform will therefore drift unless it uses energy to maintain its position [24, 52]. Movement also changes the applicable maritime regime and may conflict with data-localization obligations [49, 50], which makes position a compliance variable as well as a physical one. The connection mode carries a comparable double character, since a cable-connected facility inherits the exposure of subsea infrastructure [54] while a satellite-connected platform is vulnerable to jamming and service dependent [26, 48].

The environmental case is another important concern. Thermal discharge, radiated noise and colonization have been studied for other marine industries and never for a computing facility [22, 131]. Direct evidence is also missing for system performance. Current deployments usually report capacity and efficiency, but not achieved availability or delivered computation. Additionally, the legal status of a self-propelled platform operating on the high seas remains unclear.

Two developments could change this assessment. Publishing operational data would enable quantitative economic analysis, while formulating the joint allocation problem would provide a common object of study. Oceanic computing subsystems already have dedicated hardware and established research. The integrated system, however, has neither. This survey describes that system, summarizes the requirements disclosed over two decades of patents, and highlights what remains unmeasured.

References

[1] Y. Ichinose, M. Hayashi, S. Nomura, B. Moser, and K. Hiekata, "Sustainable data centers in southeast asia: Offshore, nearshore, and onshore systems for integrated data and power," *Sustainable Cities and Society*, vol. 81, p. 103867, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2210670722001949

[2] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, no. 6481, pp. 984–986, 2020.

[3] International Energy Agency, "Energy and AI," International Energy Agency, Paris, Special Report, 2025. [Online]. Available: https://www.iea.org/reports/energy-and-ai

[4] A. Kumar and N. M. Pindoriya, "Toward sustainable data center operation: A review on existing infrastructures, integrated smart energy management frameworks, and future perspectives," *Renewable and Sustainable Energy Reviews*, vol. 230, p. 116664, 2026. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032125013371

[5] D. Mytton, "Data centre water consumption," *npj Clean Water*, vol. 4, p. 11, 2021.

[6] J. Shuja, A. Gani, S. Shamshirband, R. W. Ahmad, and K. Bilal, "Sustainable cloud data centers: A survey of enabling techniques and technologies," *Renewable and Sustainable Energy Reviews*, vol. 62, pp. 195–214, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S136403211630051X

[7] Z. Cao, X. Zhou, H. Hu, Z. Wang, and Y. Wen, "Toward a systematic survey for carbon neutral data centers," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 2, pp. 895–936, 2022.

[8] A. A. Alkrush, M. S. Salem, O. Abdelrehim, and A. Hegazi, "Data centers cooling: A critical review of techniques, challenges, and energy saving solutions," *International Journal of Refrigeration*, vol. 160, pp. 246–262, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0140700724000458

[9] S. Cai and Z. Gou, "Towards energy-efficient data centers: A comprehensive review of passive and active cooling strategies," *Energy and Built Environment*, vol. 7, no. 1, pp. 206–226, 2026. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666123324000916

[10] S. Ölz and M. Beerepoot, "Deploying renewables in southeast asia: Trends and potentials," OECD Publishing, Paris, IEA Energy Papers 2010/06, 2010. [Online]. Available: https://doi.org/10.1787/5kmd4xs1jtmr-en

[11] L. Sirkis, "Offshore data centers powered by ocean renewable energy," National Renewable Energy Laboratory, Research Concept Proposal TN 265486, Aug. 2025. [Online]. Available: https://efiling.energy.ca.gov/GetDocument.aspx?DocumentContentId=102339&tn=265486

[12] Keppel Data Centres, "Floating data centre," 2026. [Online]. Available: https://www.keppeldatacentres.com/innovations/floating-data-centre/

[13] J. Clidaras, D. W. Stiver, and W. Hamburgen, "Water-based data center," U.S. Patent and Trademark Office, U.S. Patent 7,525,207 B2, Apr. 2009. [Online]. Available: https://patents.google.com/patent/US7525207B2/en

[14] Microsoft Research, "Project natick phase 2: Northern isles deployment," Project report, 2020. [Online]. Available: https://natick.research.microsoft.com/

[15] N. J. Madusanka, Y. Fan, S. Yang, and X. Xiang, "Digital twin in the maritime domain: A review and emerging trends," *Journal of Marine Science and Engineering*, vol. 11, no. 5, p. 1021, 2023.

[16] J.-K. Paik, "Special issue: From the titanic era to the ai era: A century of transformation in ships and offshore structures," *Ships and Offshore Structures*, vol. 21, no. 1, pp. 1–3, 2026. [Online]. Available: https://doi.org/10.1080/17445302.2025.2603413

[17] I. L. Thorson, G. A. Sheldon-Coulson, B. L. Moffat, and D. W. Place, "Reservoir-regulating digital load control," U.S. Patent and Trademark Office, U.S. Patent 11,898,530 B2, 2024. [Online]. Available: https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/11898530

[18] A. Radovanović *et al.*, "Carbon-aware computing for datacenters," *IEEE Transactions on Power Systems*, vol. 38, no. 2, pp. 1270–1280, 2023.

[19] Blue-Cloud Consortium, "Blue-cloud: Piloting innovative services for marine research and the blue economy," EU Horizon 2020 project, Grant Agreement No. 862409, 2023. [Online]. Available: https://blue-cloud.org/

[20] A. A. Periola, A. A. Alonge, and K. A. Ogudo, "Ocean warming events resilience capability in underwater computing platforms," *Scientific Reports*, vol. 14, p. 3781, Feb. 2024. [Online]. Available: https://doi.org/10.1038/s41598-024-54050-8

[21] F. S. Alqurashi, A. Trichili, N. Saeed, B. S. Ooi, and M.-S. Alouini, "Maritime communications: A survey on enabling technologies, opportunities, and challenges," *IEEE Internet of Things Journal*, vol. 10, no. 4, pp. 3525–3547, 2023.

[22] P. Vuong, A. McKinley, and P. Kaur, "Understanding biofouling and contaminant accretion on submerged marine structures," *npj Materials Degradation*, vol. 7, no. 1, p. 50, 2023.

[23] A. B. Suryawanshi, S. Borde, N. Kale, and S. Rangdale, "A comprehensive review of underwater data centers: Challenges, technologies, and future directions," in *2025 2nd International Conference on Integration of Computational Intelligent System (ICICIS)*, 2025, pp. 1–6.

[24] G. A. Sheldon-Coulson and B. L. Moffat, "Self-powered computing buoy," United States Patent and Trademark Office, U.S. Patent US 10,668,990 B2, Jun. 2020. [Online]. Available: https://patents.google.com/patent/US10668990B2/en

[25] H. Yang, H. Dong, and J. Zhu, "Underwater data centres chart a path to zero-carbon operation," *Nature Reviews Electrical Engineering*, vol. 3, pp. 76–78, 2026. [Online]. Available: https://doi.org/10.1038/s44287-025-00246-7

[26] A. Abdullah, D. Blow, S. Rampazzi, and M. J. Islam, "Active localization of close-range adversarial acoustic sources for underwater data center surveillance," *IEEE Journal of Oceanic Engineering*, 2026, early access.

[27] A. Aili, J. Choi, Y. S. Ong, and Y. Wen, "The development of carbon-neutral data centres in space," *Nature Electronics*, vol. 8, no. 11, pp. 1016–1026, 2025.

[28] W. Lin, J. Lin, Z. Peng, H. Huang, W. Lin, and K. Li, "A systematic review of green-aware management techniques for sustainable data center," *Sustainable Computing: Informatics and Systems*, vol. 42, p. 100989, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2210537924000349

[29] International Energy Agency, "Key questions on energy and AI," International Energy Agency, Paris, Tech. Rep., 2026. [Online]. Available: https://www.iea.org/reports/key-questions-on-energy-and-ai

[30] P. G. Finn, "Water-based computing system," Patent US 9,648,777 B2, 2017, priority and filing date 6 January 2014; published as US 2015/0194813 A1; granted 9 May 2017.

[31] ——, "Water-based computing system," Patent US 9,839,160 B2, 2017, continuation of US 9,648,777 B2; priority date 6 January 2014; filed 21 February 2017; published as US 2017/0164521 A1; granted 5 December 2017.

[32] E. C. Peterson, B. F. Cutler, J. Whitaker, Norman Ashton, P. Johnson, A. J. Fleming *et al.*, "Underwater container cooling via integrated and external heat exchangers," U.S. Patent and Trademark Office, U.S. Patents 9,801,313 B2 and 9,844,167 B2, 2017, applications 14/752,669 and 14/752,676, filed 26 June 2015. [Online]. Available: https://patents.google.com/patent/US9801313B2/en

[33] B. Cutler *et al.*, "Intrusion detection for submerged datacenters," U.S. Patent and Trademark Office, U.S. Patent 9,913,409 B2, 2018. [Online]. Available: https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/9913409

[34] ——, "Artificial reef datacenter," U.S. Patent and Trademark Office, U.S. Patent 10,524,395 B2, Dec. 2019. [Online]. Available: https://patents.google.com/patent/US10524395B2/en

[35] G. A. Sheldon-Coulson, B. L. Moffat, and D. W. Place, "Self-powered, self-propelled computer grid with loop topology," Patent US 2019/0 219 026 A1, 2019, priority date 17 January 2018; filed 15 January 2019; published 18 July 2019. Application reported as abandoned.

[36] B. L. Moffat and G. A. Sheldon-Coulson, "Self-powered computing buoy (continuation)," U.S. Patent and Trademark Office, U.S. Patent 11,846,264 B2, Dec. 2023. [Online]. Available: https://patents.google.com/patent/US11846264B2/en

[37] N. Kosugi, S. Takahashi, and K. Kamei, "Underwater data center," Patent US 10,136,562 B2, 2018, priority JP 2016-159497 dated 16 August 2016; US filing 3 July 2017; published as US 2018/0054916 A1 on 22 February 2018; granted 20 November 2018.

[38] T. H. Han, H. M. Hong, and J. K. Park, "Underwater data center," Patent US 12,317,453 B2, 2025, priority KR 10-2022-0127266 dated 5 October 2022; US filing 20 December 2022; published as US 2024/0121917 A1 on 11 April 2024; granted 27 May 2025.

[39] T. H. Wee and S. Kumar, "Modular floating data centre park," Patent SG 11 202 100 783T A, 2021, priority date 10 September 2018; filed 10 September 2019; published 25 February 2021; related international publication WO 2020/055330 A1.

[40] T. H. Wee, S. Kumar, T. K. Poh, and K. Zaw, "Centralized power and cooling plant with integrated cable landing station," Patent WO 2021/071 433 A1, 2021, priority date 11 October 2019; filed 9 October 2020; published 15 April 2021.

[41] V. Royer, "Floating electronic device and data center including such a device," Patent FR 3,087,083 B1, 2020, priority and filing date 9 October 2018; application published 10 April 2020; patent granted 30 October 2020.

[42] A. C. Magcale and B. Taylor, "Submerged data center facility system and method," Patent US 11,856,686 B1, 2023, priority and filing date 28 November 2018; granted 26 December 2023. The specification incorporates by reference US 9,784,460 B2, which claims the benefit of provisional application 61/861,197 filed 1 August 2013.

[43] D. Pu, R. Fu, C. Li, S. Hu, Z. He, T. Xu, N. Jiang, and W. Shen, "Underwater data center," Patent CN 216,783,820 U, 2022.

[44] United States Congress, "Patents," Title 35, United States Code, §§ 102, 103, and 112, 1952, as amended. [Online]. Available: https://www.uspto.gov/web/offices/pac/mpep/consolidated_laws.pdf

[45] European Patent Office, "Convention on the grant of european patents (european patent convention)," Articles 54, 56, and 83; consolidated text, 2000, convention originally signed in 1973; text revised by the Act of 29 November 2000 and subsequently amended. [Online]. Available: https://www.epo.org/en/legal/epc

[46] B. Ott, P. M. Wenzel, and P. Radgen, "Analysis of cooling technologies in the data center sector on the basis of patent applications," *Energies*, vol. 17, no. 14, p. 3615, 2024.

[47] J. Sheldon, W. Zhu, A. Abdullah, K. Butler, M. J. Islam, and S. Rampazzi, "Deep note: Can acoustic interference damage the availability of hard disk storage in underwater data centers?" in *Proc. 15th ACM Workshop on Hot Topics in Storage and File Systems*, 2023, pp. 51–57.

[48] J. Sheldon *et al.*, "AquaSonic: Acoustic manipulation of underwater data center operations and resource management," in *Proc. IEEE Symposium on Security and Privacy*, 2024, pp. 331–349.

[49] United Nations, "United nations convention on the law of the sea," Montego Bay, 10 December 1982; United Nations Treaty Series, vol. 1833, p. 3; entered into force 16 November 1994, 1982. [Online]. Available: https://www.un.org/depts/los/convention_agreements/texts/unclos/unclos_e.pdf

[50] European Union, "Regulation (EU) 2016/679 of the european parliament and of the council on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (general data protection regulation)," Official Journal of the European Union, L 119, 4 May 2016, pp. 1–88, 2016. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2016/679/oj

[51] United States Congress, "Clarifying lawful overseas use of data (CLOUD) act," Pub. L. No. 115-141, div. V, 132 Stat. 1213; codified in part at 18 U.S.C. § 2713, 2018.

[52] Panthalassa, "Panthalassa: Go where the energy is," Company website, 2026, accessed 5 August 2026. [Online]. Available: https://panthalassa.com/

[53] Z. Skidmore, "Panthalassa unveils wave-powered floating data center platform," Data Center Dynamics, May 2026. [Online]. Available: https://www.datacenterdynamics.com/en/news/panthalassa-unveils-wave-powered-floating-data-center-platform/

[54] M. Abner and S. Bauk, "On underwater data centers: Surveillance, monitoring, and environmental management in the Baltic sea," *Journal of Information Technology and Applications*, vol. 14, no. 2, pp. 142–149, 2024.

[55] Lin-gang Special Area Administration, "World's first underwater data center drives sustainable computing," https://en.lingang.gov.cn/html/website/lg/English/CityofInnovation/InnovationProjects/2021462605065908225.html, 2026.

[56] H. Bailey, B. Robertson, and B. Buckham, "Variability and stochastic simulation of power from wave energy converter arrays," *Renewable Energy*, vol. 115, pp. 721–733, 2018.

[57] M. Göteman, J. Engström, M. Eriksson, and J. Isberg, "Optimizing wave energy parks with over 1000 interacting point-absorbers using an approximate analytical method," *International Journal of Marine Energy*, vol. 10, pp. 113–126, 2015.

[58] H. Fang and M. Duan, *Chapter 1 - The Environment and Environmental Load of Offshore Oil Engineering*. Boston: Gulf Professional Publishing, 2014, pp. 1–140. [Online]. Available: https://www.sciencedirect.com/science/article/pii/B9780123969774000019

[59] T. Vervaet, V. Stratigaki, B. De Backer, K. Stockman, M. Vantorre, and P. Troch, "Experimental modelling of point-absorber wave energy converter arrays: A comprehensive review, identification of research gaps and design of the WECfarm setup," *Journal of Marine Science and Engineering*, vol. 10, no. 8, p. 1062, 2022.

[60] D. Pelosi, F. Gallorini, G. Alessandri, and L. Barelli, "A hybrid energy storage system integrated with a wave energy converter: Data-driven stochastic power management for output power smoothing," *Energies*, vol. 17, no. 5, p. 1167, 2024.

[61] G. Rajapakse, S. Jayasinghe, and A. Fleming, "Power smoothing and energy storage sizing of vented oscillating water column wave energy converter arrays," *Energies*, vol. 13, no. 5, p. 1278, 2020.

[62] P. Roy, H. B. Karayaka, J. He, and Y.-H. Yu, "Cost investigation of battery-supercapacitor hybrid energy storage system for grid-connected hourly dispatching wave energy converter power," *Journal of Energy Storage*, vol. 72, p. 108347, 2023.

[63] A. Allegue-Osset, C. Fernández-Blanco, L. Castro-Santos, and A. Filgueira-Vizoso, "Lcoe analysis of north sea offshore wind farms using scenario-based simulation: Weather windows and maintenance strategies," *Ocean Engineering*, vol. 363, p. 126467, 2026. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0029801826023012

[64] DNV, "Marine operations and marine warranty," DNV, Standard DNV-ST-N001, 2021. [Online]. Available: https://www.dnv.com/oilgas/download/dnv-st-n001-marine-operations-and-marine-warranty/

[65] Z. Song, A. Marburg, and D. Manalang, "Resident subsea robotic systems: A review," *Marine Technology Society Journal*, vol. 54, no. 5, pp. 21–31, 2020.

[66] N. Palomeras, A. Penalver, M. Massot-Campos, P. L. Negre, J. J. Fernandez, P. Ridao, P. J. Sanz, and G. Oliver-Codina, "I-AUV docking and panel intervention at sea," *Sensors*, vol. 16, no. 10, p. 1673, 2016.

[67] F. Hong, I. Sarantopoulos, E. Hogg, D. Richardson, Y. Zhang, H. Williams, D. Sweeney, A. Chatzieleftheriou, and A. Rowstron, "Self-maintaining [networked] systems: the rise of datacenter robotics!" in *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks (HotNets '24)*, Irvine, CA, USA, 2024, pp. 1–8.

[68] "Automated maintenance of datacenter computers using mobile robotic manipulators," Patent US 10,935,980 B2, 2021.

[69] I. Davidson, P. Cahill, A. Hinz, D. Kluza, C. Scianni, and E. Georgiades, “A review of biofouling of ships’ internal seawater systems,” *Frontiers in Marine Science*, vol. 8, p. 761531, 2021.

[70] H. Chen and D. Li, “Current status and challenges for liquid-cooled data centers,” *Frontiers in Energy Research*, vol. Volume 10 - 2022, 2022. [Online]. Available: https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2022.952680

[71] S. Ma, Y. C. Chou, H. Zhao, L. Chen, X. Ma, and J. Liu, “Network characteristics of LEO satellite constellations: A Starlink-based measurement from end users,” in *IEEE INFOCOM 2023 – IEEE Conference on Computer Communications*, 2023, pp. 1–10.

[72] J. Pan, J. Zhao, and L. Cai, “Measuring the satellite links of a LEO network,” in *ICC 2024 – IEEE International Conference on Communications*, 2024, pp. 4439–4444.

[73] L. Izhikevich, M. Tran, K. Izhikevich, G. Akiwate, and Z. Durumeric, “Democratizing LEO satellite network measurement,” *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 8, no. 1, pp. 1–26, 2024. [Online]. Available: https://doi.org/10.1145/3639039

[74] C. M. Schieler, B. C. Bilyeu, J. S. Chang, A. S. Garg, A. J. Horvath, K. M. Riesing, B. S. Robinson, J. P. Wang, S. Piazzolla, and B. Keer, “Recent on-orbit results and ARQ performance analysis for the TBIRD 200-gbps mission,” in *2023 IEEE International Conference on Space Optical Systems and Applications (ICSOS)*, 2023, pp. 49–55.

[75] N. Witternigg, M. Schönhuber, E. Leitgeb, and T. Plank, “Feasibility assessment of optical technologies for reliable high capacity feeder links,” *Acta Astronautica*, vol. 89, pp. 254–260, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0094576513001197

[76] Y. Kaymak, R. Rojas-Cessa, J. Feng, N. Ansari, M. Zhou, and T. Zhang, “A survey on acquisition, tracking, and pointing mechanisms for mobile free-space optical communications,” *IEEE Communications Surveys & Tutorials*, vol. 20, no. 2, pp. 1104–1123, 2018.

[77] D. He, Q. Wang, X. Liu, Z. Song, J. Zhou, Z. Wang, C. Gao, T. Zhang, X. Qi, Y. Tan, G. Ren, B. Qi, J. Ren, Y. Cao, and Y. Huang, “Shipborne acquisition, tracking, and pointing experimental verifications towards satellite-to-sea laser communication,” *Applied Sciences*, vol. 9, no. 18, p. 3940, 2019. [Online]. Available: https://doi.org/10.3390/app9183940

[78] European Space Agency, “World-first gigabit-per-second laser link between aircraft and geostationary satellite,” ESA news release, ScyLight programme, Feb. 2026, 26 February 2026. [Online]. Available: https://www.esa.int/Applications/Connectivity_and_Secure_Communications/World-first_gigabit-per-second_laser_link_between_aircraft_and_geostationary_satellite

[79] H. Kaushal and G. Kaddoum, “Optical communication in space: Challenges and mitigation techniques,” *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 57–96, 2017.

[80] A. F. G. Júnior, C. P. Azzolin, C. A. R. Castillo, and V. G. A. Carneiro, “Availability analysis of a ship-to-ground FSO link,” *Journal of Optical Communications and Networking*, vol. 14, no. 5, pp. 339–350, 2022.

[81] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, “Delay-tolerant networking architecture,” Internet Research Task Force, Tech. Rep. RFC 4838, Apr. 2007.

[82] L. Zago, A. N. Simos, A. Kawano, and A. M. Kogishi, “A new vessel motion based method for parametric estimation of the waves encountered by the ship in a seaway,” *Applied Ocean Research*, vol. 134, p. 103499, 2023.

[83] Y. Wang, M. N. Soltani, and D. M. A. Hussain, “Attitude stabilization of marine satellite tracking antenna using model predictive control,” *IFAC Journal of Systems and Control*, vol. 17, p. 100173, 2021.

[84] A. Lebkowski and J. Wnorowski, “A comparative analysis of energy consumption by conventional and anchor-based dynamic positioning of ship,” *Energies*, vol. 14, no. 3, p. 524, 2021. [Online]. Available: https://doi.org/10.3390/en14030524

[85] M. Nakamura, W. Koterayama, H. Kajiwara, and T. Mitamura, “Application of a dynamic positioning system to a moored floating platform,” in *Proceedings of the Fourth International Offshore and Polar Engineering Conference*. International Society of Offshore and Polar Engineers, Apr. 1994, pp. ISOPE–I–94–026. [Online]. Available: https://onepetro.org/ISOPEIOPEC/proceedings-abstract/ISOPE94/All-ISOPE94/ISOPE-I-94-026/22859

[86] Y. Gu, Z. Chuang, A. Zhang, A. Hu, and S. Ji, “Dynamic response analysis and positioning performance evaluation of an arctic floating platform based on the mooring-assisted dynamic positioning system,” *Journal of Marine Science and Engineering*, vol. 11, no. 3, p. 486, 2023. [Online]. Available: https://doi.org/10.3390/jmse11030486

[87] D. Qiao, H. Li, H. Liang, B. Wang, W. Zhang, and J. Ou, “Evaluation of the positioning capability of a novel concept design of dynamic positioning-assisted mooring systems for floating offshore wind turbines,” *Ocean Engineering*, vol. 299, p. 117286, 2024. [Online]. Available: https://doi.org/10.1016/j.oceaneng.2024.117286

[88] J. Yu, X. Cheng, Y. Fan, X. Ni, Y. Chen, and Y. Ye, “Mooring design of offshore aquaculture platform and its dynamic performance,” *Ocean Engineering*, vol. 275, p. 114146, 2023. [Online]. Available: https://doi.org/10.1016/j.oceaneng.2023.114146

[89] O. Alagili, M. A. I. Khan, S. Ahmed, S. Imtiaz, H. Zaman, and M. Islam, “An energy-efficient dynamic positioning controller for high sea conditions,” *Applied Ocean Research*, vol. 129, p. 103331, 2022.

[90] Q. Ding, F. Deng, S. Zhang, Z. Du, and H. Yang, “Multi-objective optimization for thrust allocation of dynamic positioning ship,” *Journal of Marine Science and Engineering*, vol. 12, no. 7, p. 1118, 2024.

[91] Nautilus Data Technologies, “EcoCore: Data center liquid cooling built for the future of AI,” Company website, 2026, accessed 5 August 2026. [Online]. Available: https://nautilusdt.com/ecocore/

[92] Mitsui O.S.K. Lines, Ltd., Hitachi, Ltd., and Hitachi Systems, Ltd., “MOL and hitachi launch initiative to convert used ships into floating data centers: Providing digital infrastructure to meet growing AI demand,” Press release, Mar. 2026, 30 March 2026. [Online]. Available: https://www.hitachi.com/en/press/articles/2026/03/0330/

[93] Aikido Technologies, “Aikido announces AI-grade floating datacenter technology,” Company news release, 2026, accessed 5 August 2026. [Online]. Available: https://www.aikidotechnologies.com/news-stories/aikido-announces-first-ai-grade-floating-datacenter

[94] Z. Zhou *et al.*, “Carbon-aware online control of geo-distributed cloud services,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 27, no. 9, pp. 2506–2519, 2016.

[95] M. Previsic, A. Karthikeyan, and D. Lyzenga, “In-ocean validation of a deterministic sea wave prediction (DSWP) system leveraging X-band radar to enable optimal control in wave energy conversion systems,” *Applied Ocean Research*, vol. 114, p. 102784, 2021.

[96] F. Fusco and J. V. Ringwood, “A study of the prediction requirements in real-time control of wave energy converters,” *IEEE Transactions on Sustainable Energy*, vol. 3, no. 1, pp. 176–184, 2012.

[97] ——, “Short-term wave forecasting for real-time control of wave energy converters,” *IEEE Transactions on Sustainable Energy*, vol. 1, no. 2, pp. 99–106, 2010.

[98] N. Booij, R. C. Ris, and L. H. Holthuijsen, “A third-generation wave model for coastal regions: 1. model description and validation,” *Journal of Geophysical Research: Oceans*, vol. 104, no. C4, pp. 7649–7666, 1999.

[99] M. T. Takci, M. Qadrdan, J. Summers, and J. Gustafsson, “Data centres as a source of flexibility for power systems,” *Energy Reports*, vol. 13, pp. 3661–3671, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352484725001623

[100] G. Moura Paredes, A. Tokat, and T. Thiringer, “Pre-design selection of the rated power of a heaving point absorber wave energy converter,” *Oceans*, vol. 6, no. 1, p. 13, 2025. [Online]. Available:

https://www.mdpi.com/2673-1924/6/1/13
[101] B. Guo, T. Wang, S. Jin, S. Duan, K. Yang, and Y. Zhao, "A review of point absorber wave energy converters," *Journal of Marine Science and Engineering*, vol. 10, no. 10, p. 1534, 2022. [Online]. Available: https://www.mdpi.com/2077-1312/10/10/1534
[102] J. Sjolte, G. Tjensvoll, and M. Molinas, "Power collection from wave energy farms," *Applied Sciences*, vol. 3, no. 2, pp. 420–436, 2013. [Online]. Available: https://www.mdpi.com/2076-3417/3/2/420
[103] European Parliament and Council of the European Union, "Directive (EU) 2023/959 of the european parliament and of the council of 10 may 2023 amending directive 2003/87/ec establishing a system for greenhouse gas emission allowance trading within the union," 2023, official Journal of the European Union, L 130, 16 May 2023. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2023/959/oj/eng
[104] International Maritime Organization, "IMO approves net-zero regulations for global shipping," Apr. 2025, marine Environment Protection Committee, MEPC 83. [Online]. Available: https://www.imo.org/en/mediacentre/pressbriefings/pages/imo-approves-netzero-regulations.aspx
[105] Linux Kernel Developers, "CPU performance scaling," Linux Kernel Documentation, 2026, accessed: August 2026. [Online]. Available: https://www.kernel.org/doc/html/latest/admin-guide/pm/cpufreq.html
[106] T. H. Davenport and J. G. Harris, *Competing on Analytics: The New Science of Winning*. Boston, MA: Harvard Business School Press, 2007.
[107] M. Eelsalu, R. D. Montoya, D. Aramburo, A. F. Osorio, and T. Soomere, "Spatial and temporal variability of wave energy resource in the eastern pacific from panama to the drake passage," *Renewable Energy*, vol. 224, p. 120180, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0960148124002453
[108] G. Li, G. Weiss, M. Mueller, S. Townley, and M. R. Belmont, "Wave energy converter control by wave prediction and dynamic programming," *Renewable Energy*, vol. 48, pp. 392–403, 2012.
[109] A. G. Oude Elferink, "Artificial islands, installations, and structures," in *Max Planck Encyclopedia of Public International Law*, A. Peters, Ed. Oxford University Press, Jul. 2024. [Online]. Available: https://opil.ouplaw.com/display/10.1093/law:epil/9780199231690/law-9780199231690-e247
[110] Standing Committee of the National People's Congress of the People's Republic of China, "Personal information protection law of the people's republic of china," 2021. [Online]. Available: https://personalinformationprotectionlaw.com
[111] ——, "Cybersecurity law of the people's republic of china," 2025. [Online]. Available: https://digichina.stanford.edu/work/translation-cybersecurity-law-of-the-peoples-republic-of-china-effective-june-1-2017/
[112] ——, "Data security law of the people's republic of china," 2021. [Online]. Available: https://en.spp.gov.cn/2021-06/10/c_948426_2.htm
[113] National Assembly of Vietnam, "Cybersecurity law no. 116/2025/qh15," Dec. 2025. [Online]. Available: https://english.luatvietnam.vn/lawoncybersecurityno116-2025-qh15dateddecember102025ofthenationalassembly-422396-doc1.html
[114] Personal Data Protection Commission Singapore, "Data protection obligations," Personal Data Protection Commission, 2023, see Transfer Limitation Obligation. [Online]. Available: https://www.pdpc.gov.sg/overview-of-pdpa/the-legislation/personal-data-protection-act/data-protection-obligations
[115] Personal Information Protection Commission of Japan, "Guidelines on the act on the protection of personal information: Provision to a third party in a foreign country," Personal Information Protection Commission, 2016. [Online]. Available: https://www.ppc.go.jp/personalinfo/legal/guidelines_offshore/
[116] Office of the Privacy Commissioner for Personal Data, Hong Kong, "Guidance on recommended model contractual clauses for cross-border transfer of personal data," May 2022. [Online]. Available: https://www.pcpd.org.hk/english/resources_centre/publications/files/guidance_model_contractual_clauses.pdf
[117] Government of India, "Digital personal data protection act, 2023," 2023. [Online]. Available: https://www.indiacode.nic.in/handle/123456789/22037
[118] A. Chander and U. P. Lê, "Data nationalism," *Emory Law Journal*, vol. 64, no. 3, pp. 677–739, 2015. [Online]. Available: https://scholarlycommons.law.emory.edu/elj/vol64/iss3/2/
[119] Court of Justice of the European Union, "Data protection commissioner v. Facebook Ireland Ltd and Maximillian Schrems," Case C-311/18, Judgment of the Court (Grand Chamber), 16 July 2020, ECLI:EU:C:2020:559, 2020. [Online]. Available: https://curia.europa.eu/juris/liste.jsf?num=C-311/18
[120] F. Casalini, J. López González, and T. Nemoto, "Mapping commonalities in regulatory approaches to cross-border data transfers," OECD Publishing, Paris, Tech. Rep. 248, 2021.
[121] J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten, "Lest we remember: Cold boot attacks on encryption keys," in *17th USENIX Security Symposium (USENIX Security '08)*. USENIX Association, 2008, pp. 45–60.
[122] H. Birkholz, D. Thaler, M. Richardson, N. Smith, and W. Pan, "Remote attestation procedures (rats) architecture," Internet Engineering Task Force, Tech. Rep. RFC 9334, Jan. 2023. [Online]. Available: https://www.rfc-editor.org/info/rfc9334/
[123] European Parliament and Council of the European Union, "Regulation (eu) 2016/679 (general data protection regulation)," 2016, oJ L 119, 4.5.2016, pp. 1–88. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2016/679/oj
[124] W. Zhou, Y. Yang, W. Han, C. Song, Z. Meng, H. Zhang, and Q. Han, "Effects of thermal discharge on the structuring of macrobenthic biological traits around the Sanmen nuclear power plant in the East China sea," *Marine Pollution Bulletin*, vol. 213, p. 117705, 2025. [Online]. Available: https://doi.org/10.1016/j.marpolbul.2025.117705
[125] Y. Lin, H. Cao, M. Liu, F. Qi, S. Zhang, and J. Xu, "The influence of thermal discharge from power plants on the biogeochemical environment of Daya bay during high primary productivity seasons," *Marine Pollution Bulletin*, vol. 211, p. 117408, 2025. [Online]. Available: https://doi.org/10.1016/j.marpolbul.2024.117408
[126] Q. Leng, F. Mohamat-Yusuff, K. N. Mohamed, N. S. Zainordin, and M. Z. Hassan, "Impacts of thermal and cold discharge from power plants on marine benthos and its mitigation measures: A systematic review," *Frontiers in Marine Science*, vol. 11, p. 1465289, 2024. [Online]. Available: https://doi.org/10.3389/fmars.2024.1465289
[127] L. Botero-Bolívar, O. A. Marino Sánchez, M. de Frutos, and E. Ferrer, "Assessing the underwater impact of aerodynamic noise from offshore wind turbines," *Wind Energy*, 2026, advance online publication. [Online]. Available: https://doi.org/10.1002/we.70133
[128] OSPAR Commission, "Overview of the impacts of anthropogenic underwater sound in the marine environment," OSPAR Commission, London, Biodiversity Series 441/2009, 2009. [Online]. Available: https://qsr2010.ospar.org/media/assessments/p00441_Noise_background_document.pdf
[129] International Maritime Organization, "Revised guidelines for the reduction of underwater radiated noise from shipping to address adverse impacts on marine life," MEPC.1/Circ.906/Rev.1, Nov. 2024, 15 November 2024; effective 1 December 2024. [Online]. Available: https://wwwcdn.imo.org/localresources/en/MediaCentre/Documents/MEPC.1-Circ.906-Rev.1%20-%20Revised%20Guidelines%20For%20The%20Reduction%20Of%20Underwater%20Radiated.pdf
[130] ——, "Action plan for the reduction of underwater noise from commercial shipping," Approved by the Marine Environment Protection Committee at its eighty-second session (MEPC 82), 2024, 30 September–4 October 2024.

[131] T. A. Smith and J. Rigby, "Underwater radiated noise from marine vessels: A review of noise reduction methods and technology," *Ocean Engineering*, vol. 266, p. 112863, 2022. [Online]. Available: https://doi.org/10.1016/j.oceaneng.2022.112863

[132] F. J. Rodrigo Saura, J. Ramis-Soriano, P. Poveda-Martínez, J. Carbajo, and E. Segovia, "Analysis of the underwater radiated noise generated by hull vibrations of the ships," *Sensors*, vol. 23, no. 2, p. 1035, 2023. [Online]. Available: https://doi.org/10.3390/s23021035

[133] B. Taormina, J. Bald, A. Want, G. Thouzeau, M. Lejart, N. Desroy, and A. Carlier, "A review of potential impacts of submarine power cables on the marine environment: Knowledge gaps, recommendations and future directions," *Renewable and Sustainable Energy Reviews*, vol. 96, pp. 380–391, 2018.

[134] L. Albert, F. Deschamps, A. Jolivet, F. Olivier, L. Chauvaud, and S. Chauvaud, "A current synthesis on the effects of electric and magnetic fields emitted by submarine power cables on invertebrates," *Marine Environmental Research*, vol. 159, p. 104958, 2020.

[135] OSPAR Commission, "Background document on potential problems associated with power cables other than those for oil and gas activities," OSPAR Commission, Tech. Rep. 370/2008, 2008. [Online]. Available: https://www.ospar.org/documents?v=7128

[136] A. M. Bell, M. von der Au, J. Regnery, M. Schmid, B. Meermann, G. Reifferscheid, T. Ternes, and S. Buchinger, "Does galvanic cathodic protection by aluminum anodes impact marine organisms?" *Environmental Sciences Europe*, vol. 32, p. 157, 2020.

[137] U. Gupta, Y. G. Kim, S. Lee, J. Tse, H.-H. S. Lee, G.-Y. Wei, D. Brooks, and C.-J. Wu, "Chasing carbon: The elusive environmental footprint of computing," in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2021, pp. 854–867.

[138] S. Ahmad, Z. H. Rizvi, and F. Wuttke, "Modeling thermal behavior of buried subsea power cables: A review of deterministic, data-driven, and hybrid approaches," *Frontiers in Built Environment*, vol. 12, 2026.

[139] T. Kirchgeorg, I. Weinberg, M. Hörnig, R. Baier, M. J. Schmid, and B. Brockmeyer, "Emissions from corrosion protection systems of offshore wind farms: Evaluation of the potential impact on the marine environment," *Marine Pollution Bulletin*, vol. 136, pp. 257–268, 2018.

[140] C. J. Emeana, T. J. Hughes, J. K. Dix, T. M. Gernon, T. J. Henstock, C. E. L. Thompson, and J. A. Pilgrim, "The thermal regime around buried submarine high-voltage cables," *Geophysical Journal International*, vol. 206, no. 2, pp. 1051–1064, 2016.