Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Beyond PUE: A Boundary-Aware Primary-Energy Framework for Evaluating Self- Powered Computing Infrastructure, with Application to a Hybrid Offshore Floating Data Center

Leon Sandler

sandler.leon@gmail.com

Independent Researcher

Research Article

Keywords: data center sustainability, power usage effectiveness, primary energy efficiency, self-powered computing infrastructure, carbon capture and storage, offshore floating infrastructure, sustainability metrics

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10949065/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Power usage effectiveness (PUE) is the dominant data-center sustainability metric, but it is defined only over the facility boundary: cooling and auxiliary overhead relative to the IT load. For self-powered infrastructure generating its own electricity on site, this boundary excludes exactly the losses—fuel-to-electricity conversion and any carbon-capture parasitic load—that determine whether total energy use is actually efficient. In this study, the hypothesis that facility-level PUE can systematically overestimate self-powered infrastructure sustainability is tested, and a boundary-aware primary-energy framework ($\eta_{primary} = \eta_{cycle} \times (1 - \text{CCS penalty})/\text{PUE}$) is developed to clarify that gap. The framework is applied to a proposal-stage architecture, the Hybrid Seabed-Energy Floating Data Center (HSE-FDC), offshore gas turbines, a wave-energy-converter array, carbon capture, and passive marine cooling, which had a PUE of approximately 1.04 with negligible grid dependency. Using literature-grounded parameters, the case study yields a plausible facility PUE of 1.08–1.09, which is comparable to that of real benchmarks, whereas the wave array contributes under 1.2% of the average power, which is far short of a genuine hybrid baseload. Critically, this favorable facility PUE corresponds to a primary energy efficiency of only 25.5–38.7%, below a conventional grid-tied facility (43.0%) and dramatically below a renewable-grid-tied facility (93.5%): a favorable facility PUE does not, by itself, indicate favorable total energy efficiency once onsite generation is included. The architecture's individual technologies are each already commercialized elsewhere; this study's contribution is the general framework and the demonstration that it changes the sustainability conclusion for a real architecture, generalizing to other off-grid contexts.

1. Introduction

1.1 The blind spot in facility-level sustainability metrics

Power usage effectiveness (PUE)—the ratio of total facility power to IT equipment power—is the standard metric for data-center energy efficiency and has led to substantial real improvement: conventional air/liquid-cooled facilities typically report a PUE of 1.30–1.50, whereas modern best-practice designs reach 1.10–1.20. PUE is defined strictly over the facility boundary, which is appropriate when electricity is imported from an external grid, since generation-side losses are then the grid operator's concern, not the facility's. This boundary becomes a blind spot; however, for infrastructure that generates its own electricity on site, fuel-to-electricity conversion losses and any parasitic load from emissions-control equipment occur upstream of the PUE accounting boundary and are consequently invisible to a PUE-only sustainability comparison. As self-powered and off-grid computing infrastructure—driven by grid interconnection delays—becomes more common, this blind spot is becoming increasingly consequential.

1.2 Hypothesis and framework

This study proposes and tests the hypothesis that facility-level PUE can systematically overestimate the sustainability of self-powered computing infrastructure and develops a boundary-aware primary-energy

accounting framework to make the resulting gap explicit and quantifiable (Section 3.4). The framework is deliberately general—it requires only a generation-cycle efficiency and any parasitic control-equipment load alongside the conventional PUE calculation—and is demonstrated here through application to a specific, previously unevaluated case study.

1.3 Case study: existing prior art

The case study architecture is a proposal-stage concept, the Hybrid Seabed-Energy Floating Data Center (HSE-FDC), which combines offshore associated-gas turbines, a wave-energy-converter (WEC) array, carbon capture and storage (CCS), and passive marine cooling on a floating platform. Three bodies of prior art are directly relevant and must be distinguished from any new contribution claimed here.

Marine-cooled data centers are an operationally proven concept, not a hypothesis. Microsoft's Project Natick deployed an 864-server sealed pod 36 m below the sea surface off Orkney, Scotland, for two years, reporting a measured PUE of 1.07, zero water consumption, and a failure rate one-eighth that of a matched land-based control group [7]. Natick was powered by the local renewable (wind/tidal) grid, not by onsite combustion. Separately, commercial-scale underwater facilities are now in deployment off Hainan and Shanghai, China, with a 24 MW wind-and-seawater-cooled pod project targeting PUE below 1.15 [7].

Floating offshore data center platforms colocated with offshore energy generation are also already in active commercial development and are not a novel deployment paradigm. Aikido Technologies' AO60DC integrates a 15–18 MW offshore wind turbine, battery storage, and 10–12 MW of computation on a single floating structure, with a projected PUE below 1.08; a proof-of-concept is under development in Norway, with a first commercial UK deployment targeted at approximately 2028 [8].

Gas-to-compute conversion of stranded or flared associated gas is likewise an existing, commercialized practice, not a proposed innovation. Crusoe Energy operates a commercial "flare-gas-to-compute" business that converts otherwise flared associated gas into on-site electricity for modular data centers, and xAI's Colossus facility in Memphis, Tennessee, is powered in part by on-site natural gas turbines (a deployment that has also drawn regulatory scrutiny over turbine permitting) [9]. The underlying gas-to-compute conversion concept is also the subject of granted U.S. patents covering on-site power generation from flare, stranded, and associated gas for distributed computing loads [10].

1.4 What remains genuinely open

Given this prior art, the individual components of the HSE-FDC case study—marine cooling, floating offshore siting, and gas-to-compute conversion—are each independently established or in active commercial development. What has not been quantitatively evaluated in the literature is either (a) their specific combination or (b) whether the facility-level PUE claimed for such architectures (here, approximately 1.04) reflects their true total energy efficiency once onsite generation is accounted for. This study supplies both a power balance evaluation of the specific combination and a general primary-energy framework that is applicable beyond this case study.

1.5 Research objectives

- Does facility-level PUE alone provide a reliable sustainability signal for self-powered computing infrastructure, or does it systematically exclude a materially large share of total energy loss?
- For the HSE-FDC case study specifically, what is the actual contribution of its wave-energy-converter array to total power supply, at literature-reported (not idealized) capacity factors, and what facility-level PUE is achievable relative to the real Natick and projected Aikido benchmarks?
- Accounting for gas-turbine cycle efficiency and CCS parasitic load via the proposed framework, what is the case study's total primary (fuel-to-compute) energy efficiency, and how does this compare to a conventional grid-tied and a renewable-grid-tied facility?

2. System architecture

The evaluated architecture follows the proposed-stage concept: seabed-associated gas-fired turbines provide baseload power; a leg-mounted array of point-absorber wave-energy converters provides auxiliary power; closed-loop deep-seawater heat exchange provides passive cooling; and combustion exhaust is routed through postcombustion CCS for reinjection into depleted seabed formations. Two gas-turbine configurations are evaluated, since offshore practices differ materially from onshore default: (a) simple (open) cycle, the historically dominant offshore choice due to lower weight and footprint [4, 5], and (b) combined cycle, which is materially more efficient but rarely deployed offshore—the Appomattox platform (Gulf of Mexico, operational since 2019, four 27 MW aeroderivative gas turbines plus a 40 MW steam turbine, ~ 50% net efficiency) is, to the authors' knowledge, the primary offshore precedent for this configuration [4, 5].

3. Mathematical model and methodology

As in the accompanying analysis to this manuscript, every number in Section 4 *is traced to a specific equation and is cited below; parameters not available from the established literature (e.g., specific platform CAPEX) are explicitly flagged as assumptions rather than asserted as results.*

3.1 Gas turbine baseload

Two scenarios are modeled using the Appomattox unit sizing as a representative reference point [4, 5]: simple-cycle, 4 × 27 MW aeroderivative units at 33% net efficiency (midpoint of the 30–35% typical range for offshore simple-cycle aeroderivative turbines); and combined-cycle, the same 4 × 27 MW units plus a 40 MW steam turbine at 50% net efficiency, matching the reported Appomattox figure.

3.2 Wave energy converter array

The wave power flux per unit crest width follows the standard deep-water formulation [1]:

$$P/w = (\rho\, g^2/64\pi) \times H_s^2 \times T_e$$

where ρ = 1025 kg/m³, g = 9.81 m/s², H_s is the significant wave height, and T_e is the energy period. A representative mid-latitude offshore wave climate (H_s = 2.5 m, T_e = 9 s, which is consistent with the 5.8–28.5 kW/m resource range reported across the assessed WEC deployment sites [3]) yields a flux of 27.6 kW/m. The array-rated outputs use a representative capture width of 20 m per buoy, 12 buoys, and 75% mechanical-to-electrical power-take-off efficiency [2]. Critically, the rated (nameplate) output is not the same as the average delivered output: real deployed point-absorber devices report capacity factors of approximately 12–51% of the rated output depending on the device design and site (e.g., Wavebob 40–51% at 12–21 kW/m reference; RM3 16% at 34 kW/m; AquaBuOY ~ 12% at 25 kW/m) [3]. A representative 25% capacity factor is applied here, avoiding the continuous/near-rated output implied in the original proposal-stage concept.

3.3 Facility PUE

Facility overhead (distinct from gas-turbine generation efficiency; see Section 3.4) is modeled as the sum of seawater circulation pump power (3% of the IT load, reflecting a larger, nonsealed platform relative to Natick's compact sealed pod), power conversion/UPS losses (4%), and control/monitoring systems (1.5%), yielding PUE = 1 + 0.085 = 1.085. This is deliberately not set below the real Natick benchmark of 1.07 without justification; the platform modeled here is a larger, more mechanically complex, nonsealed structure than Natick's single sealed capsule, and 1.085 is presented as plausible and comparable to, not superior to, the best demonstrated real-world case.

3.4 Primary energy efficiency

Facility PUE alone does not capture the efficiency of onsite power generation, which a grid-tied facility does not need to account for (that loss occurs at the utility, off the facility's books) but which an off-grid, self-generating platform must do. This study therefore defines and reports a primary energy efficiency metric:

$$\eta_{primary} = \eta_{cycle} \times (1 - CCS_{penalty})/PUE$$

where η_cycle is the gas-turbine net efficiency (Section 3.1) and CCS_penalty is the fraction of gross generation consumed by carbon capture (Section 3.5). This metric is compared against a grid-tied baseline (mixed-fuel grid generation efficiency 52%, transmission/distribution efficiency 95%, facility PUE 1.15, all representative modern figures) and a renewable-grid-tied baseline consistent with Natick (no onsite combustion; renewable electricity treated as the reference case, PUE 1.07).

This comparison requires a stated convention for what counts as "primary energy" across fundamentally different generation types, and the choice materially affects the result: This study uses the direct/physical-energy-content method, in which each generation source's primary energy is its own energy content at the point of generation—one kWh of renewable electricity is counted as one kWh of primary energy, with no fossil-fuel-equivalent inflation (i.e., not the partial-substitution method some aggregate energy statistics use for renewables). This is the more conservative choice for this study's own thesis: the substitution method would count renewable primary energy as the larger fossil input it

displaces (fossil plants require approximately 2–3 kWh of fuel per kWh of electricity), which would lower, not increase, the reported efficiency of the renewable-tied comparator. The 93.5% Natick-style figure reported below is therefore a conservative (least favorable to this study's fossil-unfavorable conclusion) choice of accounting convention, not a figure selected to maximize the gap against the HSE-FDC.

3.5 CCS energy penalty and carbon balance

The postcombustion CCS parasitic energy penalty is modeled at a literature-consistent midpoint of 16%, bounded by an 11% thermodynamic lower bound and an assumed 25% upper bound for less-optimized retrofits [6]. Standard 90% CO_2 capture-rate design targets are assumed. Combustion CO_2 is computed from the standard natural gas emission factor (0.202 kg CO_2 per kWh thermal input [11]), scaled by cycle efficiency and CCS-driven additional fuel consumption, and then reduced by the assumed capture rate. A US grid average of 0.386 kg CO_2/kWh [12] is used as the grid-tied comparison baseline.

4. Results

4.1 Wave contribution to the average power supply

At the baseline wave climate and array specification (Section 3.2), the array's rated capacity is 4.97 MW; at a 25% capacity factor, the average delivered output is 1.24 MW. Compared with a 108–149 MW gas-turbine baseload (simple- and combined-cycle scenarios, respectively), this represents only 0.8–1.1% of the average total power supply (Table 1, Fig. 1)—not a meaningful “hybrid” or “auxiliary baseload” contribution, as characterized in the original proposal, but a minor supplement whose primary practical value would likely be fuel offset during favorable sea states rather than firm capacity.

Table 1
Average power contribution of the gas turbine and wave array by cycle scenario

Scenario	Gas turbine rated (MW)	Wave array avg. (MW)	Wave share of avg. supply
Simple-cycle (established offshore practice)	108.0	1.24	1.14%
Combined-cycle (Appomattox-precedented)	148.0	1.24	0.83%

4.2 Facility PUE vs. real-world benchmarks

The modeled facility PUE (1.085) lies between Microsoft's measured Natick result (1.07, renewable-grid-tied) and Aikido's projected AO60DC target (< 1.08, wind-tied, in development) and is comfortably below the 1.30–1.50 range typical of conventional terrestrial hyperscale facilities and the 1.15 midpoint of modern best-practice terrestrial designs (Fig. 2). This is a plausible, literature-consistent claim -- unlike the original proposal's ~ 1.04 figure, which was not supported by any stated calculation and which this model does not reproduce even under favorable assumptions.

4.3 Primary energy efficiency: the central finding

The primary energy efficiency metric defined in Section 3.4 is shown in Fig. 3. Despite its favorable facility PUE, the simple-cycle HSE-FDC configuration delivers only 25.5% primary energy efficiency, and the combined-cycle configuration delivers 38.7% – both below the 43.0% computed for a conventional mixed-fuel grid-tied facility and dramatically below the 93.5% figure for a renewable-grid-tied facility such as Natick. This is the central finding of this study: a favorable facility PUE does not indicate favorable total energy efficiency when the facility also performs its own fossil-fuel generation because the thermodynamic conversion loss (67% for simple-cycle, 50% for combined-cycle gas turbines) and the CCS parasitic load (16%) both occur upstream of the PUE accounting boundary and are consequently invisible to a PUE-only comparison.

4.4 Carbon intensity

The carbon intensity per kWh delivered to the IT load is shown in Fig. 4. Without CCS, both gas-turbine scenarios exceed the US grid average (simple cycle: 0.664 kg CO_2/kWh vs. grid average 0.386 kg CO_2/kWh); with CCS at the assumed 90% capture rate, the net intensity falls to 0.079–0.052 kg CO_2/kWh, which is below the grid average but still substantially above that of a renewable-grid-tied facility such as Natick (~ 0.02 kg CO_2/kWh, dominated by upstream renewable-infrastructure lifecycle emissions rather than combustion). CCS is therefore not optional for this architecture to achieve a credible carbon-intensity claim relative to grid alternatives; without it, the platform is carbon intensive relative to the grid it claims to bypass.

5. Discussion

5.1 Novelty must be reframed

Given the prior art in Section 1.2, this architecture's genuine contribution cannot rest on marine cooling (proven by Natick), floating offshore siting (in active commercial development by Aikido), or gas-to-compute conversion (commercialized by Crusoe, deployed by xAI, and patented [10]). The contributions of this study are narrower and more specific: a quantitative feasibility and primary-energy accounting of the seabed-gas-plus-wave-plus-CCS combination specifically, and the finding that this particular combination, despite being a competitive facility PUE, does not achieve competitive total primary energy efficiency relative to either grid-tied or renewable-tied alternatives.

5.2 The wave energy contribution is not material

At 0.8–1.1% of the average supply, the wave-energy-converter array does not materially change the platform's power balance, carbon intensity, or primary energy efficiency; removing it entirely from the architecture would change every result in Section 4 by less than 1.2 percentage points. Its inclusion in the original proposal as a named coequal power source alongside gas alongside wind (“auxiliary offshore wind and closed-loop marine thermal exchange”) is not supported by wave-resource physics at

representative offshore capacity factors, and future iterations of this concept should describe it as a minor supplement to fuel offset, not a hybrid baseload component.

5.3 Scope of the sustainability claim

This study's findings should not be read as a claim that the HSE-FDC case study is inherently less sustainable than all existing data-center infrastructure, nor does the proposed framework claim that self-powered architectures are inherently inferior in general. Modern data-center sustainability practices increasingly combine high-efficiency IT hardware, workload consolidation, advanced liquid and free cooling, renewable power purchase agreements, thermal-aware workload scheduling, and waste-heat reuse; this study does not model or compare these methods against any of these operational strategies and makes no claim about them. The comparison developed here is narrower and specific to the accounting boundary problem: for a facility that performs its own onsite fossil-fuel generation, facility PUE alone omits the generation-side losses that a grid-tied facility does not need to report, and this can produce a misleadingly favorable sustainability impression. The results of the case study precisely illustrate that the facility PUE of the HSE-FDC (1.08–1.09) is competitive with real, renewable-tied marine-cooled infrastructure, whereas its primary energy efficiency (25.5–38.7%) is not. This is a testable, falsifiable claim about the specific case study evaluated, not a general theorem about all self-powered or distributed computing architectures, and the original proposal's framing of “zero grid dependency” as an unqualified advantage should be understood in this narrower, corrected sense.

5.4 Broader applicability of the framework

The boundary-aware primary-energy framework developed in Section 3.4 is not specific to offshore platforms or to gas-turbine generation; it applies to any computing infrastructure that performs onsite energy conversion rather than importing already generated electricity. Candidate contexts include remote AI training clusters sited at stranded or curtailed renewable generation (where curtailment-to-compute schemes still involve a generation-side efficiency and, if paired with backup fossil generation, a conversion loss analogous to Section 4.3); microgrid-connected facilities that blend on-site generation with limited grid import, for which a blended primary-energy efficiency could be computed using the same accounting structure; and other stranded-energy computing proposals, including additional flare-gas-to-compute deployments beyond the Crusoe and xAI examples cited in Section 1.3, for which the same gap between facility PUE and primary energy efficiency should be expected wherever on-site combustion is involved. Applying the framework to these contexts is proposed as future work; this study establishes an accounting method and demonstrates it on one fully worked case study.

5.5 Reliability advantage of self-powered infrastructure

The efficiency and carbon-intensity findings above are deliberately narrow and should not obscure a separate, genuine benefit of self-generating architectures such as HSE-FDC: insulation from main-grid disruptions. Grid-tied facilities inherit the reliability of the utility grid they connect to. U.S. electricity customers experienced an average total interruption duration of about two hours per year when major events are excluded, but roughly eight hours per year once major weather and grid events are included,

with this total trending upward in years with severe storm activity [13]. Large grid-tied data-center clusters are not immune to grid-side risk either: in a widely reported July 2026 incident, a single transmission-line fault in Northern Virginia caused approximately 60 data centers to instantaneously drop roughly 3 GW of demand from the grid, destabilizing frequency across a large section of the eastern U.S. grid before operators restored balance within minutes [14]. A fully self-powered, off-grid architecture such as HSE-FDC has no equivalent exposure: because it generates all of its own power on site and carries no grid interconnection to fail, it is structurally immune to utility-side outages, transmission-line faults, and grid-stability events of the kind described above, independent of its primary energy efficiency. This reliability benefit is orthogonal to the efficiency and carbon findings of this study: a low primary energy efficiency does not diminish it, and a self-powered facility with poor primary energy efficiency can still be more resilient to grid disruption than a highly efficient grid-tied one. This study does not attempt to quantify reliability (e.g., an uptime or expected-outage-hours metric) alongside efficiency and carbon intensity, and doing so is identified as future work; a complete comparison would also need to weigh this reliability benefit against the self-powered platform's own failure modes, such as offshore maintenance logistics and single-point turbine or wave-array failures with no grid to fall back on.

5.6 CCS reinjection and combined-cycle weight are unresolved engineering risks

The original proposal’s plan to compress and reinject CO_2 “into depleted seabed geological formations” assumes reinjection capacity, geological suitability, and monitoring infrastructure that this study does not evaluate and that are not trivial at gigawatt-scale continuous operation; this is a stated limitation (Section 5.8), not a resolved design element. Separately, the combined-cycle configuration, while precedented at Appomattox, adds substantial topside weight (an HRSG and steam turbine train reported at 48–179 tonnes in comparable studies [4]), which is a material engineering trade-off against a floating, rather than fixed, platform's buoyancy and stability margins—a trade-off that the original proposal's combined-cycle assumption did not address.

5.7 Regulatory precedent

xAI’s Colossus facility has drawn documented regulatory scrutiny over gas turbines, permitting an onshore, grid-adjacent site [9]; an offshore, internationally sited platform would face a distinct but not necessarily lighter regulatory regime, particularly for CCS reinjection permitting and emission monitoring, and this is not addressed by the original proposal’s risk-analysis section, which covered only corrosion, fuel treatment, and biofouling.

5.8 Limitations

- The wave-power flux formula (Section 3.2) is a standard linear-wave-theory approximation and does not capture higher-order sea-state effects, device-array hydrodynamic interactions, or mooring/structural dynamics.

- CCS reinjection logistics, capacity, and long-term seabed storage integrity are not modeled and are treated as open engineering questions, not resolved design elements.
- Combined-cycle weight/footprint trade-offs specific to a floating (as opposed to fixed) platform are noted qualitatively (Section 5.6) but not quantitatively modeled here.
- Capital cost and levelized cost of computation are not modeled in this study; the contribution here is restricted to power balance, PUE, and primary-energy accounting.
- Grid and renewable-tied comparison baselines (Section 3.4) use representative, literature-typical figures rather than a site-specific comparison, since no specific onshore alternative site is named in the original proposal.
- Embodied infrastructure impact, computing-resource utilization, and workload-management strategies—all recognized levers in the broader sustainable-computing literature—are outside the scope of this study's accounting framework, which addresses only the generation-and-cooling energy boundary; a full lifecycle comparison would need to incorporate them and is identified here as future work rather than modeled.

6. Conclusions

This study aimed to test the hypothesis that facility-level PUE can systematically overestimate the sustainability of self-powered computing infrastructure and to develop a general accounting framework capable of detecting that gap. Applied to a proposal-stage case study (HSE-FDC) that had a PUE of approximately 1.04 with negligible grid dependency, the framework reveals the following: (1) a facility PUE of 1.08–1.09 is plausible and comparable to, although not superior to, real-world marine-cooled benchmarks (Natick 1.07, Aikido's projected < 1.08); (2) at realistic point-absorber capacity factors, the case study's wave-energy array contributes under 1.2% of the average power supply and should not be characterized as a hybrid baseload component; and (3) despite its favorable facility PUE, the case study's primary (fuel-to-compute) energy efficiency of 25.5–38.7% is below both a conventional grid-tied facility (43.0%) and, by a wide margin, a renewable-grid-tied facility (93.5%) because on-site combustion and CCS parasitic load impose losses upstream of the PUE accounting boundary. This confirms the study's central hypothesis for at least this one case: a favorable facility PUE did not correspond to favorable total energy efficiency. The results are reported as a property of this specific case study and accounting boundary, not as a general claim that self-powered or distributed computing architectures are inherently less sustainable than centralized, grid-tied alternatives; modern operational sustainability strategies (workload optimization, advanced cooling, renewable procurement, heat reuse) remain outside this study's scope and are not compared against. The proposed value of the framework is broader than that of this case study: it offers a boundary-aware method applicable to other self-powered and off-grid computing contexts—remote AI clusters, microgrid-connected facilities, and stranded-energy computing—wherever facility PUE alone would otherwise be the only reported sustainability signal.

Declarations

Ethics approval and consent to participate

Not applicable. This study is a theoretical modeling and literature-based feasibility analysis; it does not involve human participants, human data, or animal subjects.

Consent for publication

Not applicable.

Competing interests

The author declares that there are no competing interests.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability

The full model implementation (gas-turbine, wave-energy-converter, PUE, primary-energy-efficiency, and carbon-intensity calculations) and all the parameters used to generate the results in this study are openly available at https://github.com/sandlerleon/hse-fdc-primary-energy-framework (MIT license) and archived with a permanent DOI on Zenodo: https://doi.org/10.5281/zenodo.22459658. The manuscript itself is additionally archived on Zenodo: https://doi.org/10.5281/zenodo.22459680.

Authors' contributions

Leon Sandler: conceptualization, methodology, software, formal analysis, investigation, writing—original draft, writing—review and editing. The author read and approved the final manuscript.

Acknowledgments

Not applicable.

References

1. Falnes J. A review of wave-energy extraction. Mar Struct. 2007;20(4):185–201. https://doi.org/10.1016/j.marstruc.2007.09.001.
2. Faizal M, Ahmed MR, Lee Y-H. A design outline for floating point absorber wave energy converters. J Renew Energy. 2014;2014:846097. https://doi.org/10.1155/2014/846097.
3. Aderinto T, Li H. Review on power performance and efficiency of wave energy converters. Energies. 2019;12(22):4329. https://doi.org/10.3390/en12224329.
4. IPIECA. (2022) Combined-cycle gas turbines. Energy Efficiency Compendium. International Petroleum Industry Environmental Conservation Association, London. https://www.ipieca.org/resources/energy-efficiency-solutions/power-and-heat-generation/combined-cycle-gas-turbines
5. Offshore Magazine. (2023) Combined cycle power plants may help offshore producers decarbonize. Endeavor Business Media, Tulsa, OK. https://www.offshore-mag.com/production/article/14294254/combined-cycle-power-plants-may-help-offshore-producers-decarbonize
6. House KZ, Harvey CF, Aziz MJ, Schrag DP. The energy penalty of post-combustion CO2 capture and storage and its implications for retrofitting the U.S. installed base. Energy Environ Sci. 2009;2(2):193–205. https://doi.org/10.1039/B811608C.
7. Microsoft. (2020) Project Natick Phase 2 results summary. Microsoft Research, Redmond, WA. https://natick.research.microsoft.com/
8. Aikido Technologies. (2026) Aikido launches offshore datacenter technology to deploy GW-scale, sovereign AI compute onboard their floating wind platform. GlobeNewswire, 3 March 2026. https://www.globenewswire.com/news-release/2026/03/03/3248183/0/en/Aikido-Launches-Offshore-Datacenter-Technology-to-Deploy-GW-scale-Sovereign-AI-Compute-Onboard-Their-Floating-Wind-Platform.html
9. Crusoe Energy Systems. (2024) How a natural gas-powered Bitcoin miner became a darling of climate tech. Crusoe Newsroom. https://www.crusoe.ai/resources/newsroom/how-a-natural-gas-powered-bitcoin-miner-became-a-darling-of-climate-tech; Data Center Dynamics (2025) Musk's xAI gets go-ahead for 41 natural gas turbines in Mississippi to power Colossus data centers. https://www.datacenterdynamics.com/en/news/musks-xai-gets-go-ahead-for-41-natural-gas-turbines-in-mississippi-to-power-colossus-data-centers/
10. Crusoe Energy Systems Inc. (2020) Systems and methods for generating and consuming power from natural gas. US Patent 10,862,307 B2, issued 8 Dec 2020.
11. US Environmental Protection Agency. Emission factors for greenhouse gas inventories. EPA; 2023. https://www.epa.gov/climateleadership/ghg-emission-factors-hub. Washington, DC.
12. US Energy Information Administration. Electric power annual 2022. Washington, DC: EIA; 2023. https://www.eia.gov/electricity/annual/.
13. US Energy Information Administration. (2021) U.S. electricity customers experienced eight hours of power interruptions in 2020. Today in Energy, EIA, Washington, DC.

https://www.eia.gov/todayinenergy/detail.php?id=50316

14. Data Center Dynamics. (2026) Virginia narrowly avoided power cuts when 60 data centers dropped off the grid at once. https://www.datacenterdynamics.com/en/news/virginia-narrowly-avoided-power-cuts-when-60-data-centers-dropped-off-the-grid-at-once/

Figures

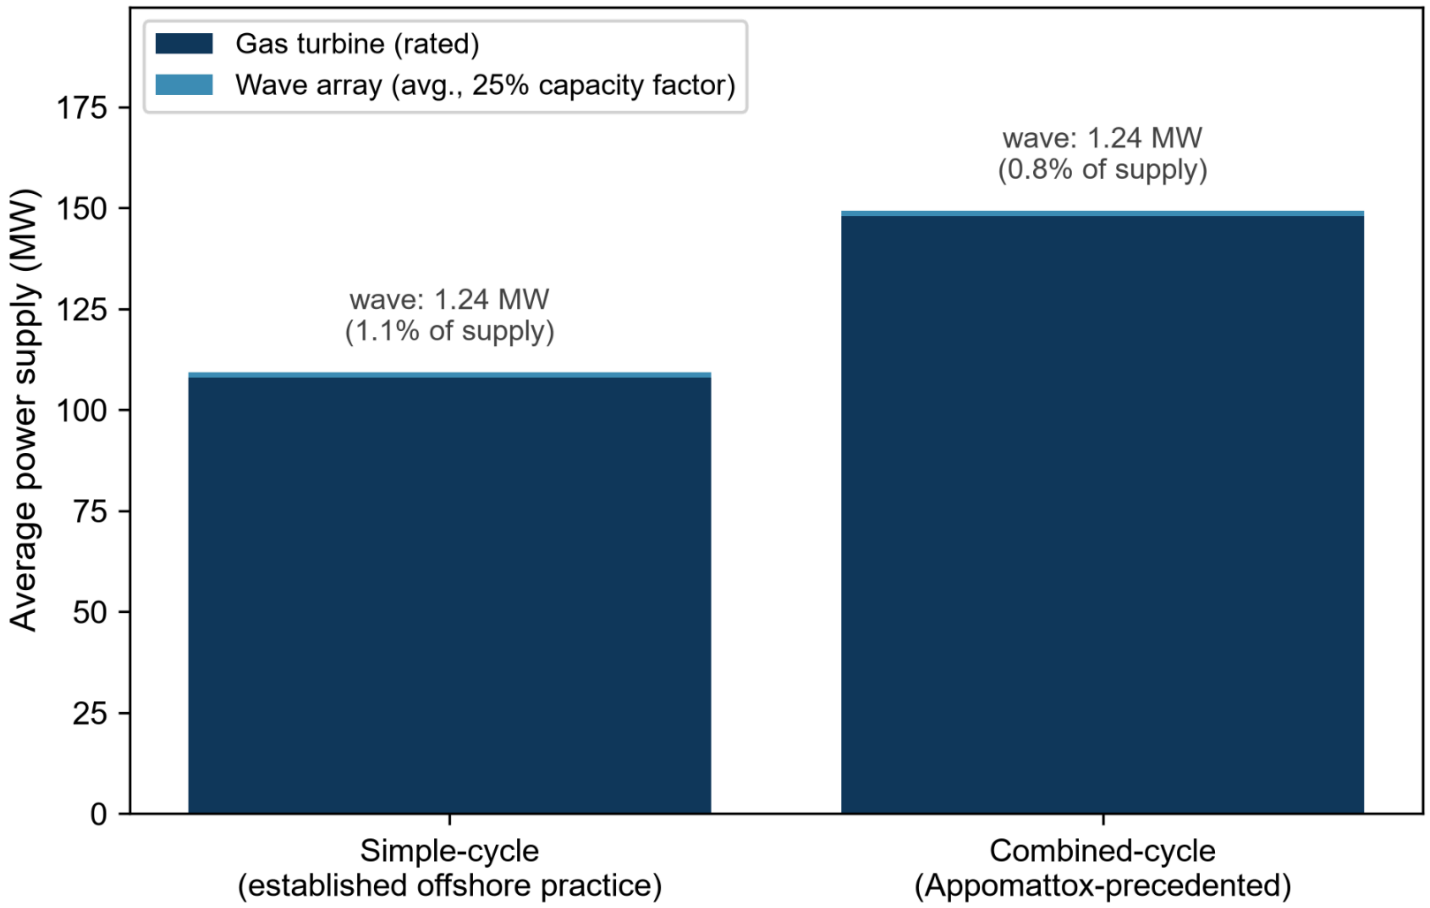


Figure 1

Contribution of the gas turbine vs. the wave array to the average power supply in both gas–turbine cycle scenarios

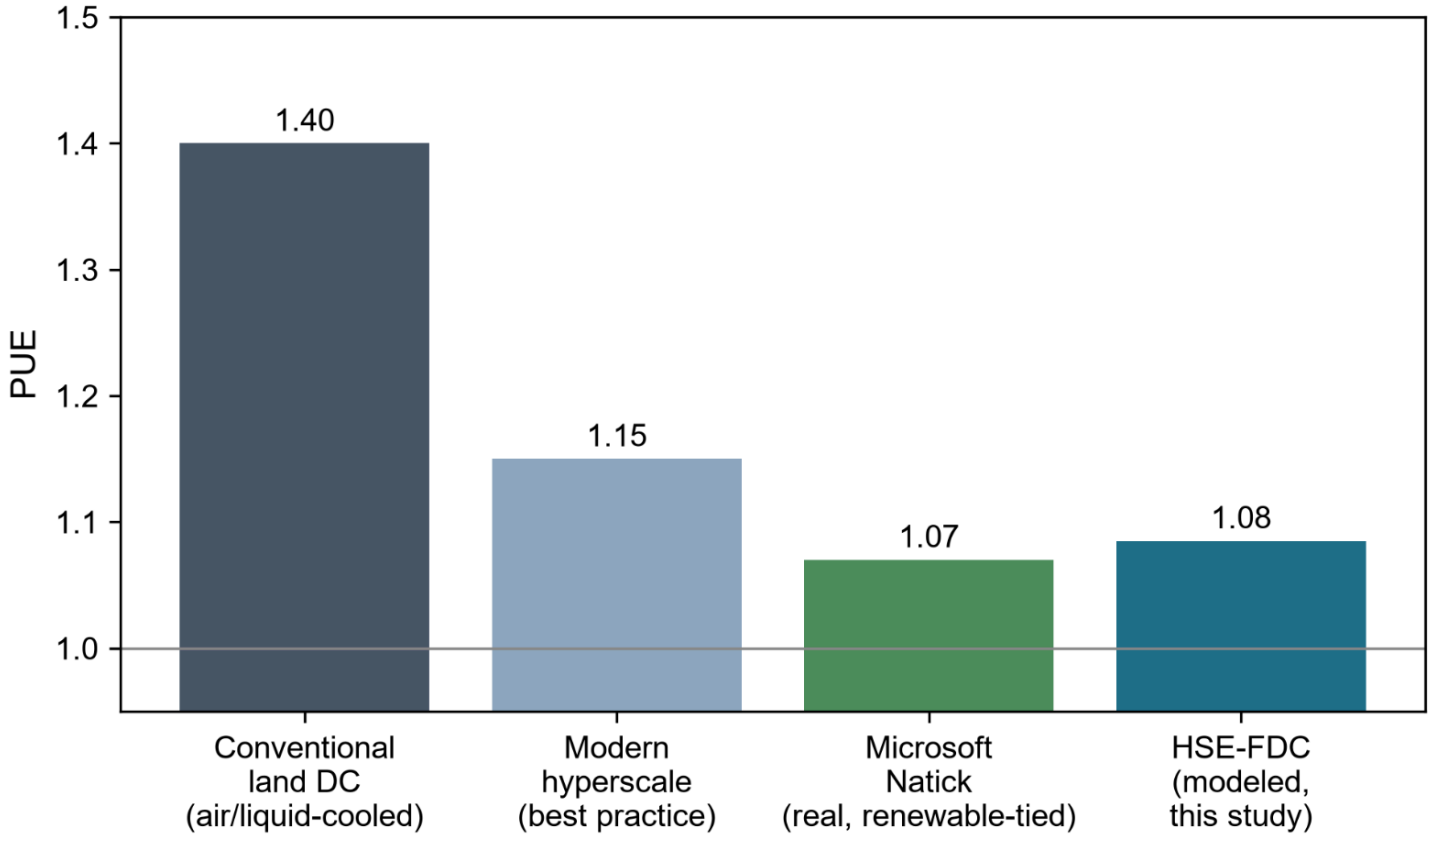


Figure 2

Facility PUE: HSE-FDC modeled value vs. conventional, best-practice, and real-world marine-cooled benchmarks

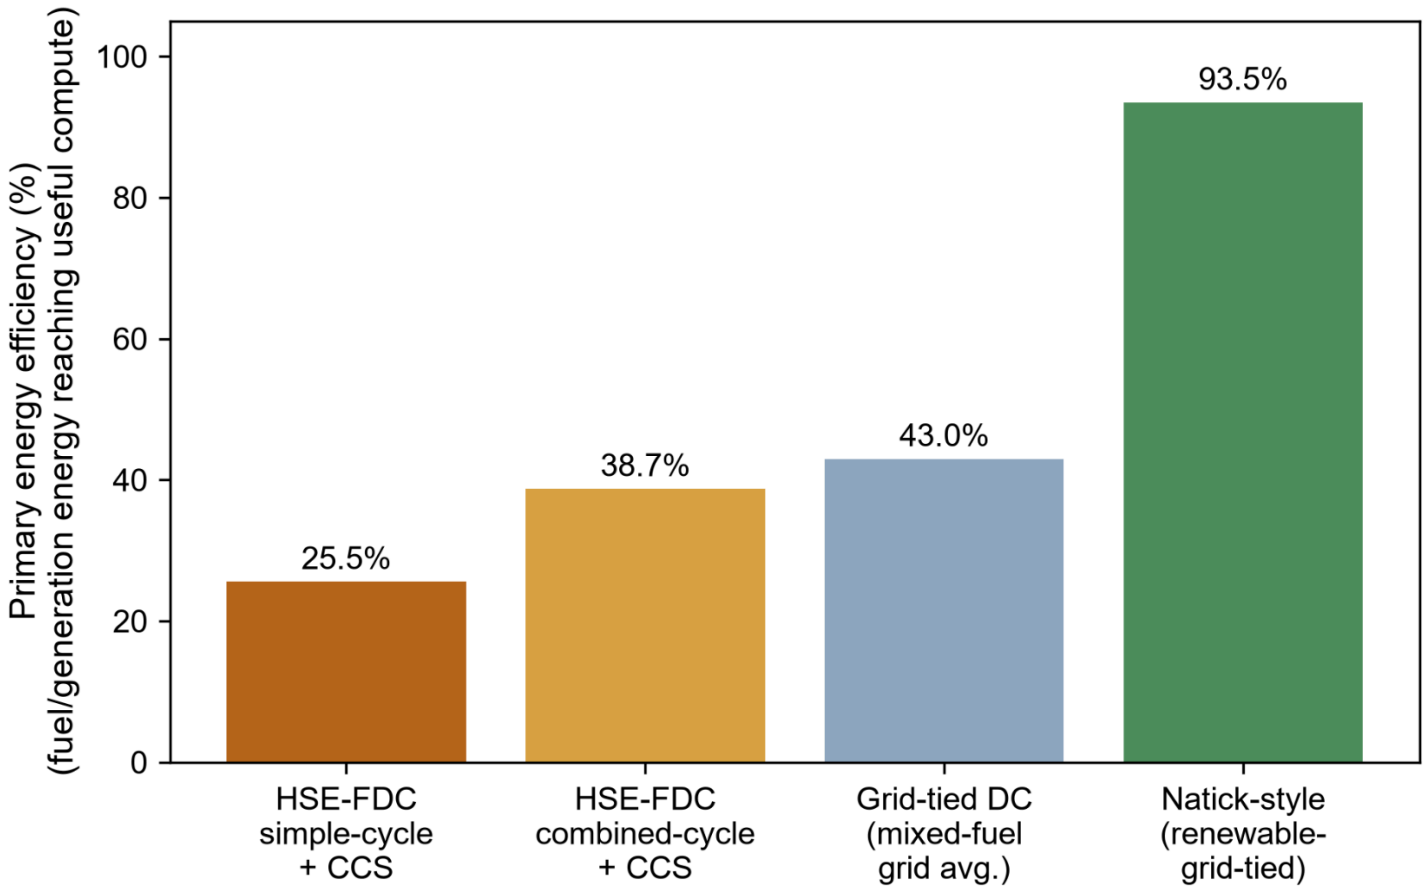


Figure 3

Primary (fuel-to-compute) energy efficiency: HSE-FDC vs. grid-tied and renewable-tied benchmarks. This metric, not the facility PUE alone, determines the total energy efficiency for a self-generating platform

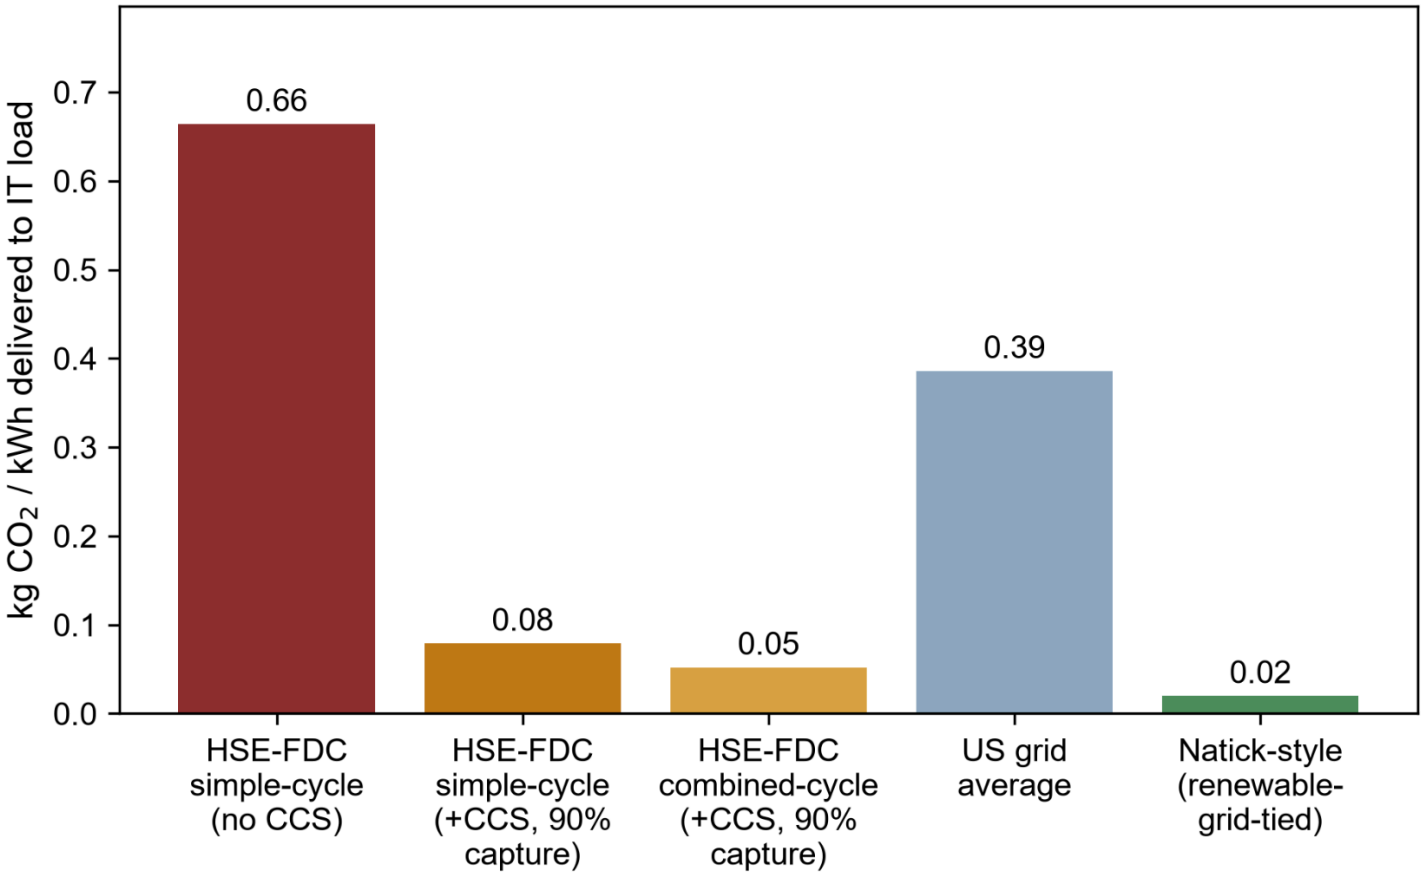


Figure 4

Carbon intensity (kg CO_2 per kWh delivered to the IT load): HSE-FDC with and without CCS, vs. grid and renewable-tied benchmarks