AN ALTERNATIVE ARCHITECTURAL FRAMEWORK FOR MOLECULAR ENERGY LOGISTICS AND NET-DECOUPLED 24/7 COMPUTE INFRASTRUCTURE

Standard Specification, Risk Management Protocol (NRG-ALADDIN)
& Global Energy Logistics White Book

Author: Ryszard Dzikowski
DESERTSOLAR UG, Berlin, Germany

Patent Application: DE 10 2026 003 523.6
German Patent and Trade Mark Office (DPMA), 22 June 2026

Revision 19

TABLE OF CONTENTS

SUPPLEMENTARY ANNEXES INCLUDED IN THIS CONSOLIDATED EDITION

These two supplementary annexes form part of this consolidated edition but retain their own internal structure, figure numbering, evidence boundaries and reference registers.

FOREWORD: THE LOGISTICS OF THE ENERGY TRANSITION

"We may take it then that an army without its baggage-train is lost; without provisions it is lost; without bases of supply it is lost."

— Sun Tzu, The Art of War, Chapter VII, trans. Lionel Giles [39]

Sun Tzu's statement describes a physical reality that reaches far beyond warfare: capability does not arise merely from the quantity of available resources, but from the ability to store them, move them across distance and time, and deliver them where and when they are required. Armies with superior numbers have failed when their supply lines broke, their provisions arrived too late, or their bases of supply were too distant from the decisive point.

The infrastructure supporting rapidly expanding segments of the global information economy faces a comparable logistical stress test in August 2026. Rapidly advancing artificial-intelligence systems are creating a growing requirement for continuous, high-density, 24/7 energy. Parallel to the long-term expansion of wind and solar generation, the need is increasing to make variable generation usable across time, distance and system boundaries. Yet an electricity system must balance supply and demand in real time. Transmission lines can move electricity, but they do not store it; interconnection capacity is finite, and generation frequently emerges far from the places where firm new loads are seeking connection.

The result is the infrastructure paradox examined in this White Book. At generation nodes, usable low-carbon energy may be curtailed, throttled, left unconnected or exposed to deeply negative prices. At load nodes, grid operators and municipalities impose connection delays, construction pauses and moratoria because transmission, distribution, water, permitting or local generation capacity cannot support the requested expansion. Hyperscalers can deploy data-centre buildings and modern accelerators far faster than conventional grid reinforcements can be planned and energised. During a five-to-eight-year interconnection delay, a substantial part of the economically relevant life cycle of the installed accelerators may pass before the compute asset receives productive power.

This White Book therefore challenges the assumption that every energy constraint must be answered primarily by constructing another power plant. It applies the logic of logistics—and, at the coordination layer, the logic of packetised digital networks—to energy availability. The relevant question is not only how much capacity exists, but how energy can be collected, buffered, certified, routed and delivered under defined service conditions.

Where electron transport through the grid is constrained, a complementary pathway is to decouple energy in time and space: use available electricity for qualified conversion or compression, place the resulting molecular energy in standardised and traceable modules, hold it as rotating inventory, and move it by road, rail, inland waterway or sea to a contracted conversion point. This does not eliminate the public grid, UHVDC, batteries, pipelines or geological storage. It creates an additional logistics layer for locations, seasons and load profiles that those systems cannot serve on the required timetable.

NRG4NOW is therefore neither a new generation technology nor a political manifesto. It is an open digital and physical assembly line for energy: a vendor-neutral architecture linking capture, compression or conversion, modular storage, transport, telemetry, safety, provenance, dispatch and destination-side power generation. Its purpose is to make existing but spatially or temporally constrained energy usable before new dedicated generation becomes the default response.

Sun Tzu identified the underlying problem faced by armies dependent on baggage trains, provisions and supply bases. The comparison used here is analytical in that it applies the principle of timely supply from military campaigns to modern energy systems. Based on this principle, the author of this White Book developed the open framework of the NRG4NOW system architecture for engineering validation and industrial implementation.

The infrastructure decisions of the AI age will not be determined by certificates and accounting claims alone. They will be determined by whether usable energy reaches the required place at the required time. The decisive routes of the AI age are therefore also routes of global energy logistics.

Ryszard Dzikowski
Berlin, August 2026

DOCUMENT STATUS AND NON-INVESTMENT NOTICE

> This White Book presents a strategic and technical system architecture for further engineering, regulatory and commercial validation. It is not an investment prospectus, financing offer or investment recommendation. Project-specific performance, cost and bankability require independent verification. Measured data, source-based observations, scenario assumptions and design targets are distinguished in the text and in the Evidence-Status Register.

HOW TO READ THIS WHITE BOOK

This White Book connects three analytical levels. They should not be read as competing products or as interchangeable claims. The first level defines the system problem, the second describes an open architecture, and the third tests carrier-specific applications and deployment corridors.

Level	Purpose	Principal content
System problem	Identify unavailable or stranded energy	Hyperscalers, grid constraints, installed versus delivered energy, curtailment and seasonal mismatch
Architecture	Create interoperability and logistics	NRG4NOW interfaces, qualified modules, hubs, orchestration, safety and governance
Applications	Test bounded use cases	Brownfields, offshore gas, China corridors, methane, hydrogen, pyrolysis and compute infrastructure

NRG4NOW is the vendor-neutral system architecture. TSTM is one optional reference concept within an open class of qualified molecular storage modules. NRG-ALADDIN is the associated risk and validation protocol; it is not a separate energy product. [1, 29]

Evidence status. Time-stamped screenshots and repeated cross-system observations are used to identify patterns and formulate testable hypotheses. They are kept distinct from measured curtailment, official redispatch records and independently validated engineering performance. Bracketed reference numbers identify the corresponding entry in the Reference Register.

This consolidated edition includes two supplementary evidence annexes: a China system analysis and a time-stamped cross-system observation study covering 3–13 August 2026.

CHAPTER 1: EXECUTIVE SUMMARY & THE INFRASTRUCTURE PARADOX

NRG4NOW is not a new generation technology and not a replacement for UHVDC or the public grid. It is a vendor-neutral logistics and interoperability layer for energy that cannot be used directly at the source or transported economically as electrons. TSTM is one optional reference concept for modular molecular storage; its performance assumptions remain subject to independent engineering review and pilot acceptance testing. [1, 29]

The proposal is not based on the assumption that maritime CNG is commercially untested. Earlier World Bank commercialization studies and the reported operation of JAYANTI BARUNA provide relevant precedents; the remaining task is to qualify an open, multi-OEM and stage-gated system architecture. [13, 14]

1.1 The AGI Compute Collision With Analog Grids

The global information economy stands at a structural breaking point in August 2026. The worldwide transition from legacy cloud workloads to continuously processing Artificial General Intelligence (AGI) and Large Language Model (LLM) training clusters demands a fundamental transformation of energy delivery systems. While conventional IT infrastructures exhibited pronounced diurnal load profiles with predictable peaks and valleys, modern AI supercomputing clusters require an uninterrupted, unyielding 24/7/365 baseload of electrical power.

This exponentially growing energy appetite collides directly with the physical, geographical, and administrative limits of electrical transmission grids globally. The power grid remains an invention rooted in the late 19th and early 20th centuries. It was conceptually engineered for synchronous, dispatchable thermal power plants (coal, natural gas, nuclear) located in close proximity to demand centers, where generation instantaneously matched consumption in real time.

The rapid, uncoordinated expansion of volatile renewable energy sources—specifically onshore/offshore wind and utility-scale solar photovoltaics—has inverted this paradigm. Because power grids transport electrons without providing inherent large-scale chemical or mechanical buffer storage at the grid level, every additional unit of unbuffered nameplate capacity installed at generation nodes introduces systemic instability.

The resulting structural phenomenon is the Infrastructure Paradox of the 2020s: A historic high in globally installed generation capacity meets systematic energy destruction through forced grid curtailment at the source, while at the consumption nodes, municipalities and grid operators enforce moratoriums, build-stops, and interconnection bans because physical grid capacity cannot support new multi-megawatt or gigawatt loads.

1.2 The Time-to-Energy and Stranded-Capital Dilemma

For leading technology enterprises (Hyperscalers)—including Amazon, Microsoft, Alphabet (Google), Meta, and Apple—this infrastructure paralysis has escalated into an existential threat to corporate valuations and Return on Investment (ROI). For fiscal year 2026, these entities have pledged cumulative capital expenditures (Capex) exceeding $660 billion, primarily earmarked for AI data center development. [6]

However, these massive capital allocations are trapped in the stranded-capital dilemma: Constructing a stateof-the-art data center building and populating it with high-density GPU accelerators currently takes 12 to 18 months using modular construction techniques. Conversely, securing an interconnection agreement and grid energization for loads exceeding 100 Megawatts (MW) in North America and Western Europe currently requires 60 to 96 months (5 to 8 years). [4, 6, 33]

This time asymmetry between compute deployment and grid delivery creates stranded capital at a multi-billion-dollar scale. Fully built compute facilities sit idle as empty shells while operators incur unearned Opex penalties and forfeit token generation revenues. The bottleneck of the AI revolution is no longer silicon availability or software optimization—it is Time-to-Energy (TTE).

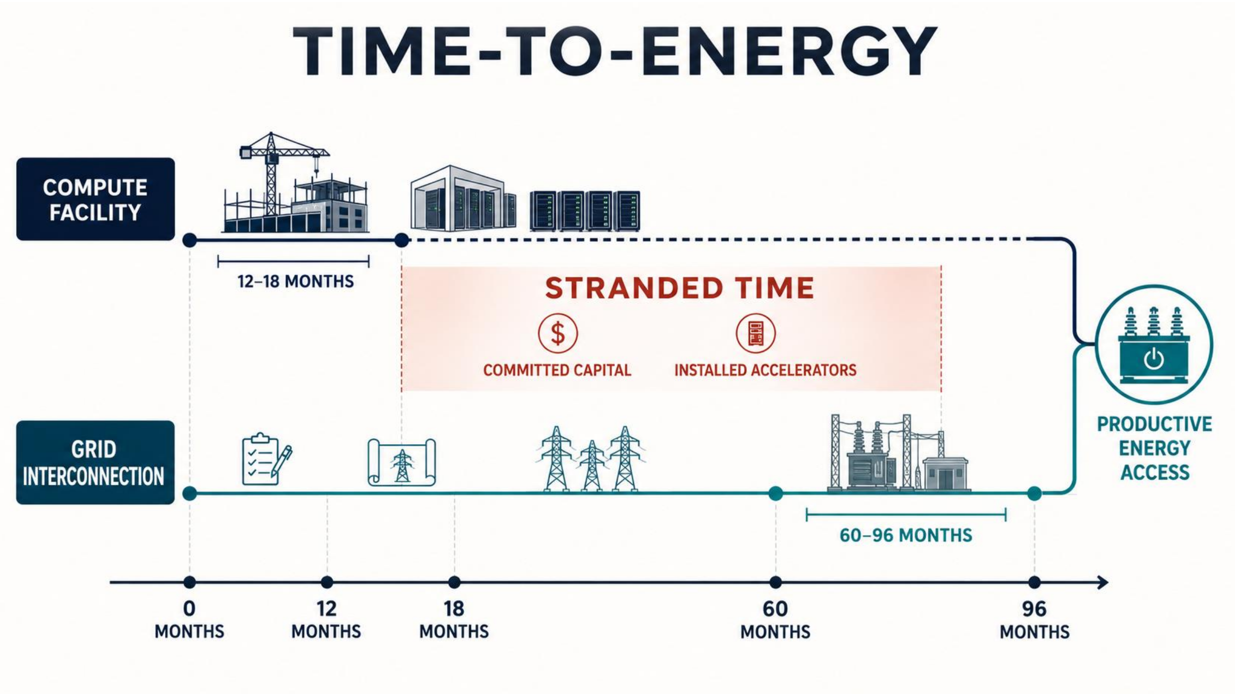

Figure 1.1: *Time-to-Energy Asymmetry – Timeline of compute facility deployment (12–18 months) versus legacy grid energization (60–96 months), illustrating the multi-year stranded capital and accelerator obsolescence interval before reaching productive energy access.*

1.3 The August 10, 2026 Catalyst: Amazon's 7.65 GW "GW Ranch" Gas Plant

The desperation among Hyperscalers was demonstrated on August 10, 2026, by the announcement of Amazon's "GW Ranch" project in Texas. To bypass multi-year interconnection queues and municipal power bans, Amazon financed the construction of an off-grid 7.65 Gigawatt (GW) natural gas power plant consisting of 35 gas turbines. [32]

This generation fleet is dedicated exclusively to powering a new AI data center cluster—completely isolated from the public transmission grid. This project illustrates two key industry realities:

1. The Forced Off-Grid Paradigm: Operators must decouple from the public grid to remain competitive.

The Logistics Trap: Lacking an open, standardized logistics layer, operators may default to new on-site generation instead of using existing energy resources. This can increase capital exposure, fossil-fuel dependence and public opposition.

The NRG4NOW framework can support a defined net-decoupled operating mode for qualified pilot sites, subject to sufficient inventory, conversion capacity, redundancy, backup and regulatory approval.

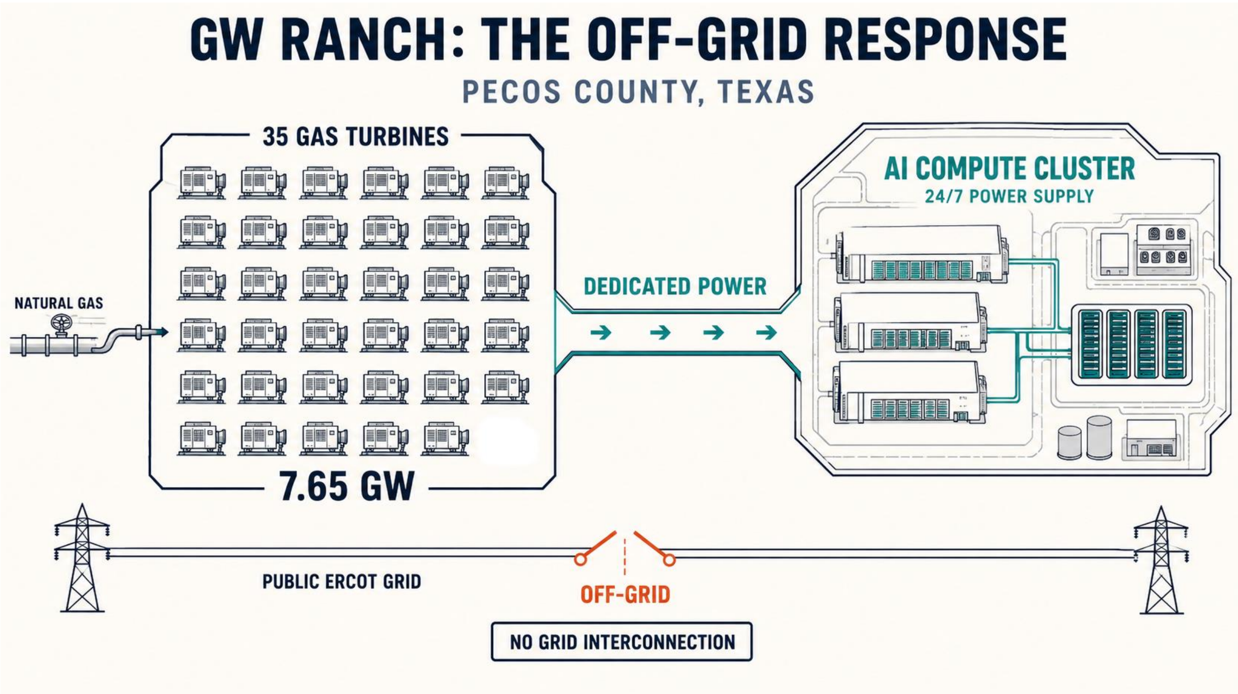

Figure 1.2: *GW Ranch (Pecos County, Texas) – The Off-Grid Response to Grid Paralysis. Dedicated 7.65 GW on-site power campus comprising 35 natural gas turbines directly powering an AI compute cluster, completely decoupled from the public ERCOT transmission network.*

1.4 The August 2026 Infrastructure Constraint: 474 GW in Texas and 500+ US Bans

During the first week of August 2026, the traditional grid-tethered data center expansion model experienced systemic failure in the United States:

On August 3, 2026, Texas Governor Greg Abbott issued an executive directive halting the processing of data center requests in the Electric Reliability Council of Texas (ERCOT) queue. Interconnection requests had reached 474 Gigawatts—over five times the historic peak demand of the entire Texas grid. ERCOT immediately suspended its Batch Zero transmission study. [4]

Governor Abbott conditioned future approvals on three audit criteria: [4]

1. Data centers must demonstrate self-provided generation (On-Site Power) rather than relying on the ERCOT grid. [4]

2. Projects must supply and recycle their own cooling water without impacting municipal water tables. [4]

3. Projects must operate without state or local tax abatements and subsidies. [4]

Simultaneously, WIRED Magazine published an investigation on August 4, 2026 (How Data Centers Broke American Politics), detailing how anti-data-center opposition has unified disparate political groups—from conservative populists to progressive environmental organizations. A Gallup poll cited in the report indicated that 70% of Americans oppose data center developments in their communities, resulting in over 345 documented project rejections or blocks in the Bryce database. [3, 5]

1.4.1 The National Siting Constraint: More Than 500 US Jurisdictions

By August 2026, disclosures attributed to The Information and the New York Post Business Report indicated that more than 500 US jurisdictions had blocked or prohibited large data-center construction by law, emergency moratorium or development ban. More than 150 municipalities and counties reportedly adopted emergency restrictions during July 2026 alone. The opposition drivers are water consumption, HVAC noise and the fear of cascading electricity-price increases for households. [6]

The Robert Bryce Data Center Rejection Database documented 345 rejected or blocked projects by the relevant August 2026 update. This journalistically compiled event database is not an official census, but it provides a traceable project-level indicator of the scale and geography of local opposition. [5]

Goldman Sachs Commodities Research and JPMorgan assessments reported that only 50% to 60% of announced data-center capacity could be energized on schedule. JPMorgan reported that almost 60% of capacity planned for 2027 had not yet entered construction despite more than $750 billion in announced 2026 AI-infrastructure investment. These figures are treated as market assessments rather than guaranteed outcomes. The grid-tethered expansion model has nevertheless become both a physical and a socio-political constraint on frontier-model deployment. [6]

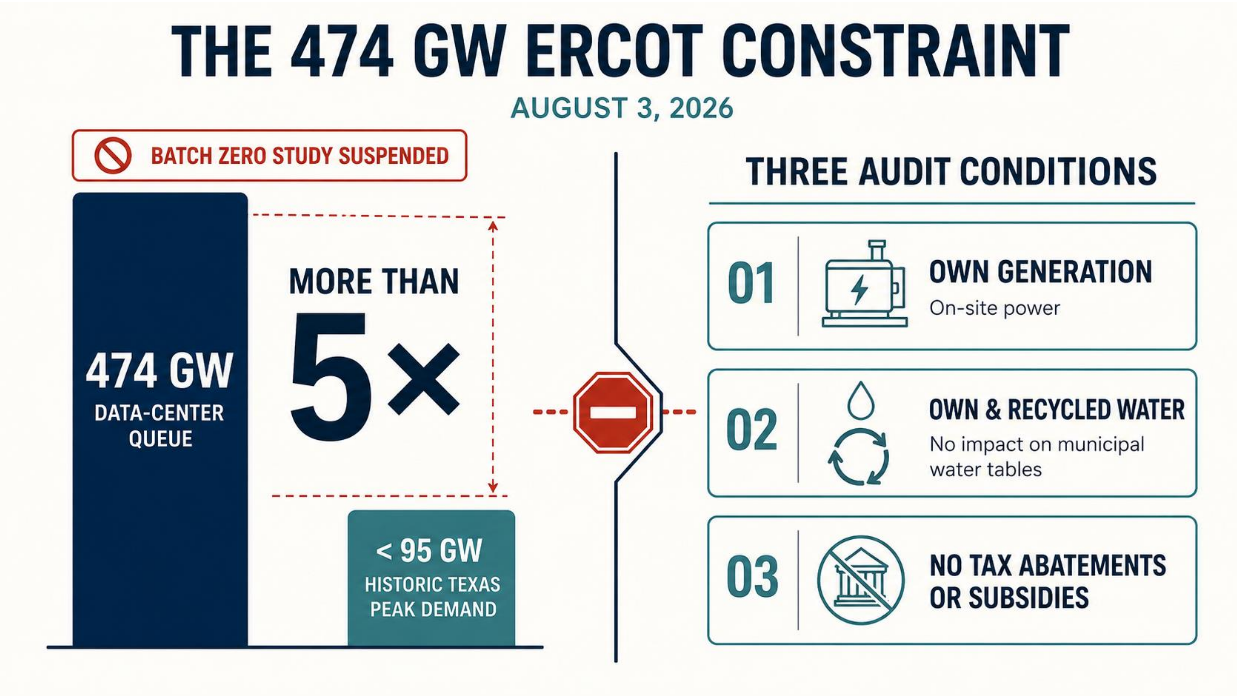

Figure 1.3: *The 474 GW ERCOT Constraint & Abbott Audit Conditions (August 3, 2026) – Structural mismatch between requested data-center load (474 GW) and historic peak grid demand (< 95 GW) leading to the suspension of ERCOT's Batch Zero study, alongside the three non-negotiable audit criteria for future site approvals.*

Economic question: Is the value of earlier available net energy greater than the full cost, conversion loss and execution risk of modular energy logistics? Chapter 7 develops this question as a time-to-energy comparison rather than assuming that transported molecular energy is automatically cheaper than a grid connection.

CHAPTER 2: SYSTEM DIAGNOSIS: THE EXPANDED GIGAWATT FALLACY

2.1 Power (GW) Versus Delivered Energy (GWh)

A core misconception in modern energy policy is the Gigawatt Fallacy—the conflation of nameplate capacity (Gigawatts, GW) with delivered energy over time (Gigawatt-hours, GWh or Terawatt-hours, TWh).

Progress in energy transitions is frequently reported in terms of installed nameplate capacity. However, nameplate capacity represents only a theoretical maximum under ideal conditions. Compute clusters and industrial facilities require delivered energy integrated over time:

A 100 MW solar PV installation in Central Europe typically yields an annual capacity factor of 10% to 12% (approx. 900 to 1,100 full-load hours). Adding more unbuffered nameplate capacity to an oversupplied grid during peak generation hours does not increase system reliability; it increases peak-to-load misalignment and forces curtailment. [8]

$$E = \int_{t_2}^{t_1} P(t) \cdot dt$$

2.2 Evidence and Energy-Availability Accounting

Curtailment, pre-emptive throttling, market non-use and non-connected capacity describe different forms of unavailable energy. They must not be added without compatible definitions, time boundaries and meter-level evidence. This White Book therefore separates official statistics, observed telemetry and explicit scenarios rather than presenting an adjusted global curtailment total as a measured fact.

2.2.1 Energy-Availability Ledger

NRG4NOW uses a source-to-load ledger so that different forms of unavailable energy are not counted twice. A project must identify the physical meter, counterfactual dispatch state, time interval and legal ownership for every claimed megawatt-hour.

Category	Accounting boundary	NRG4NOW treatment
Active curtailment	Metered potential minus instructed output at a connected plant	Eligible only when the dispatch instruction and baseline are auditable
Pre-emptive throttling	Output withheld before a formal curtailment event	Separate scenario class; requires plant-control or market evidence
Non-connected capacity	Project or plant without an energized export route	Opportunity set, not generated or curtailed energy
Market non-use	Technically available output not dispatched for economic reasons	Value depends on price, contract and alternative use

2.3 US and EU Interconnection Queues

Queue figures must be read with their reporting year and technology scope. The frequently cited United States figure of approximately 2,600 GW describes the 2025 queue dataset. Within that reporting frame, approximately 1,840 GW related to solar and wind, while the broader queue including storage and other technologies approached 3,890 GW of requested capacity. [33]

After removal of withdrawn, duplicate and otherwise non-viable entries, the 2026 working estimate used in this White Book is approximately 2,200 GW for the United States. The corresponding European

Union queue estimate is approximately 1,700 GW. The combined US–EU working quantity is therefore 3,900 GW of requested connection capacity. [17, 33]

Region	Queue quantity used	Interpretation [17, 33]
United States	2,200 GW	Adjusted 2026 working estimate
European Union	1,700 GW	Reported connection-queue estimate
Combined US + EU	3,900 GW	Requested capacity; not installed generation

Queue capacity is not installed capacity and cannot be added directly to annual curtailed energy. GW describes power; TWh describes energy. Any conversion of the 3,900 GW queue into annual TWh therefore requires an explicit capacity-factor assumption and does not imply that every queued project will be built.

2.4 Queue-Based Energy Scenarios and Hyperscaler Equivalence

For analytical purposes, this White Book converts the combined 3,900 GW US–EU interconnection queue into annual energy scenarios. At an assumed weighted average capacity factor of 32.2%, the theoretical annual output would be approximately 11,000 TWh: 3,900 GW × 8,760 hours × 0.322 ≈ 11,000 TWh per year.

This is a theoretical scenario quantity, not measured generation, measured curtailment or a forecast. Queue capacity includes different technologies, development stages and duplicate or non-viable projects, many of which may never be built.

The table below shows how the annualized energy result changes with the assumed weighted capacity factor. The 850 TWh curtailment quantity is shown separately as an additional scenario input and must not be interpreted as part of the interconnection queue.

Weighted capacity factor	Annualized output of 3,900 GW queue	Plus 850 TWh curtailment scenario
20.0%	6,833 TWh/year	7,683 TWh/year
27.5%	9,395 TWh/year	10,245 TWh/year
32.2%	11,000 TWh/year	11,850 TWh/year
35.0%	11,957 TWh/year	12,807 TWh/year

The hyperscaler comparison below uses only the explicitly identified 32.2% scenario, including 850 TWh of assumed annual curtailment. It expresses an energy ratio, not an assertion that queued projects can supply a hyperscaler continuously or that the energy is physically available at its sites. The comparison must be recalculated whenever the queue scope, realization rate, capacity factor, curtailment quantity or corporate demand estimate changes.

Hyperscaler	Estimated annual electricity demand (2026)	11,850 TWh scenario equivalent [6]
Microsoft / Azure	23.5 TWh/year	approximately 504 years
Meta Platforms	18.4 TWh/year	approximately 644 years
Amazon AWS	34.0 TWh/year	approximately 349 years
Alphabet / Google Cloud	26.0 TWh/year	approximately 456 years
Combined four hyperscalers	101.9 TWh/year	approximately 116 years

2.5 The 10-Day Empirical System Study: 3–13 August 2026

To eliminate a single-day sampling effect, the analysis covers the continuous period from 3 to 13 August 2026. Central Europe and the US Sun Belt experienced a persistent high-pressure weather pattern with predominantly cloud-free conditions and strong solar irradiance. Synchronized Electricity Maps, Windy and market telemetry exposed a structural separation between installed capacity and usable grid output. [8]

The individual Electricity Maps and Windy screenshots, timestamps and observation notes are assigned to a separate empirical data annex. The table below is the synthesis; the annex preserves the underlying visual evidence and allows every observation to be checked against its time, zone and displayed data source. [8]

Region / market	Installed PV	Observed midday utilization	Peak PV grid share / buffer [8]
Germany	~124.0 GW	41.3%–45.0%	71%–80%; limited synchronous sink
Portugal	~3.2–3.5 GW	90.0%–107.2%	~45%; pumped-hydro buffer
Greece	~9–10 GW	90.0%–98.5%	~55%; export and load
Denmark	country dataset	up to ~97%	small fleet / strong interconnection
Texas / ERCOT	30.3 GW AC listed	100.0%–103.5%	~38.5%; cooling-load match

2.5.1 Germany's 41%–45% Nameplate Dilemma

Across the observation window, the German PV fleet delivered approximately 41.37% to 45% of its listed 124 GW at the midday peak despite cloud-free conditions. On 13 August, 51.4 GW represented 41.37% of installed capacity while covering a large share of the 72.4 GW grid load. The apparent ceiling combines three effects: the rooftop mandate and east-west geometry, behind-the-meter self-consumption, and real electronic curtailment or redispatch of controllable parks. [8]

2. Rooftop geometry: more than 60% of German PV capacity is installed on homes and commercial roofs. Fixed east-west orientations broaden the daily curve but cannot reach the simultaneous noon peak represented by the aggregated kWp register. [8]
3. Invisible power: part of the generated electricity is consumed behind the meter by factories, heat pumps, households and private batteries and therefore does not appear in public-grid telemetry. [8]
4. Curtailment and redispatch: when the public grid and residual demand cannot absorb the available peak, controllable plants are reduced electronically to protect lines and voltage conditions. [8]

2.5.2 Southern European Buffering and the Texas Load Match

Portugal and Greece reached 90% to more than 100% of listed capacity during the same period. Their fleets contain a larger share of optimally oriented or tracking utility-scale plants. Portugal can route midday surplus into pumped-hydro operation, including the Tâmega complex. Texas combines DC/AC overbuild of approximately 1.4:1 with an extreme synchronous air-conditioning load: more than 81 GW of summer demand can absorb a solar peak above 31 GW. [8]

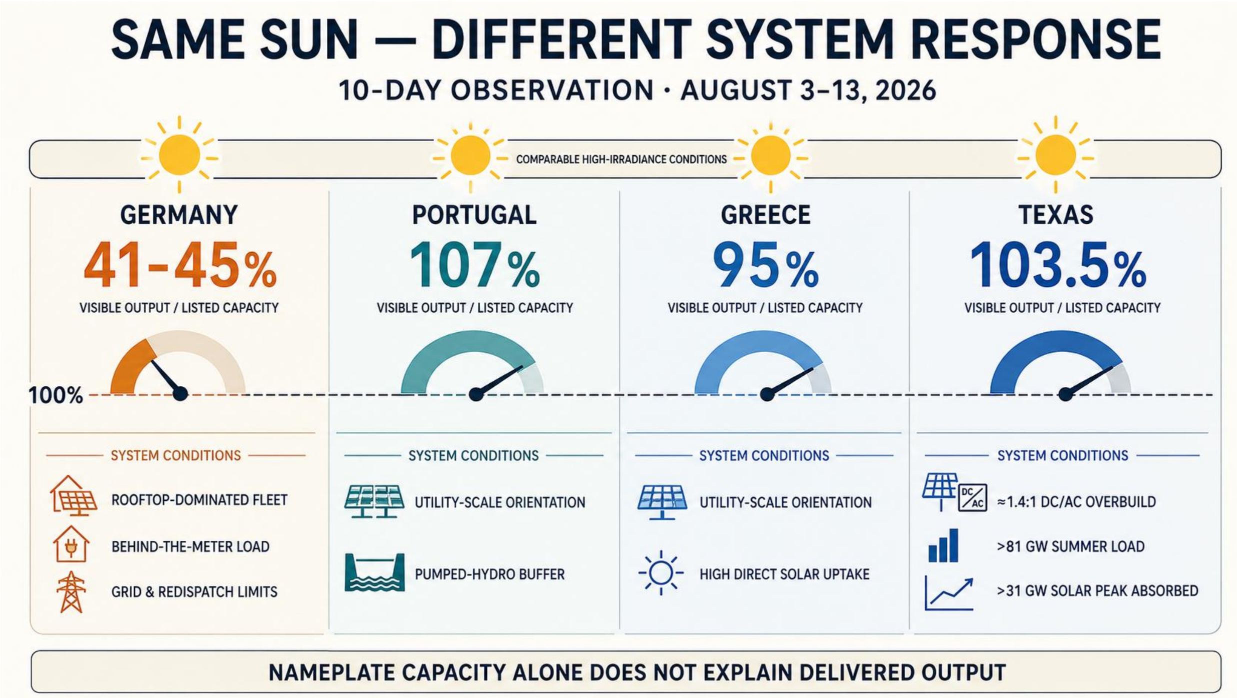

Figure 2.1: *Same Sun — Different System Response (10-Day Observation, August 3–13, 2026). Cross-system empirical comparison of visible solar output relative to listed nameplate capacity across Germany (41–45%), Portugal (107%), Greece (95%), and Texas ERCOT (103.5%), highlighting structural differences in fleet geometry, mechanical buffering, and synchronous demand absorption.*

2.5.3 Cross-System Wind Observation: Germany, West Denmark and ERCOT

The author used repeated, time-stamped Electricity Maps observations to compare visible wind-generation profiles across several systems. The method follows an investigative sequence: document recurring anomalies at different locations, formulate an explanatory hypothesis, and then test that hypothesis against meteorology, prices, grid flows and published dispatch or redispatch records. [31]

During 14–17 August 2026, Germany displayed a comparatively limited visible wind contribution while solar dominated the daytime production profile. West Denmark maintained a substantial wind contribution, including periods of simultaneous solar generation, and its negative net-flow profile showed that cross-border exports acted as a spatial buffer. ERCOT likewise maintained a substantial wind contribution alongside its daily solar peaks and large regional load. [31]

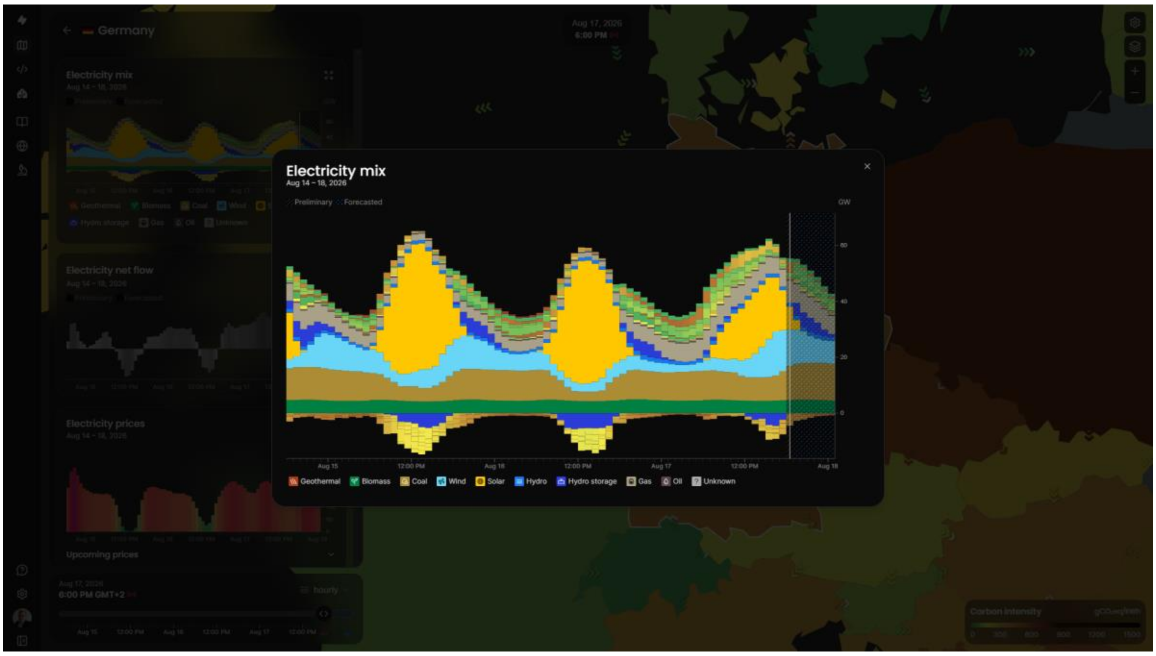

Figure 2.2a: Germany electricity mix, 14–18 August 2026 [31]

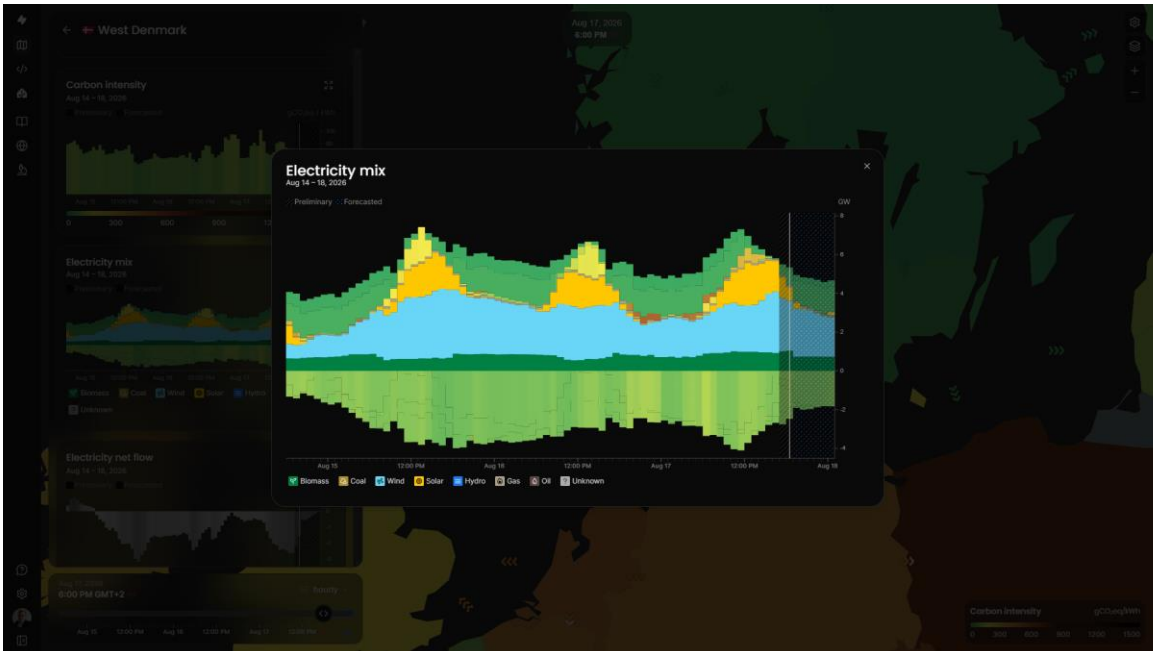

Figure 2.2b: West Denmark electricity mix and net exports, 14–18 August 2026 [31]

Figure 2.2c: ERCOT electricity mix, 14–17 August 2026 [31]

Interpretation boundary. These screenshots reveal a pattern of observation; however, they do not constitute a direct measurement of the limitation and should be regarded as hypothesis-generating observations. Aggregated generation alone cannot separate wind availability, maintenance, market dispatch, congestion and active reduction. The resulting hypothesis is therefore narrower: the wind energy admitted into a power system may depend not only on the available resource, but also on transmission capacity, export capability, storage, demand and regional dispatch practice. [31]

2.6 California: Curtailment and BESS Saturation

CAISO telemetry demonstrates the duration limit of battery energy storage. Between noon and 3:00 PM, battery charge acceptance falls as the fleet approaches 100% state of charge. Four-hour systems then discharge through the evening peak and are depleted around midnight. Natural gas and interstate imports can subsequently account for up to 70% of supply. [7]

The 4.5 TWh curtailment record in the first half of 2026 confirms that adding BESS does not itself create seasonal storage. Batteries balance hours; they do not carry August energy into January. [7]

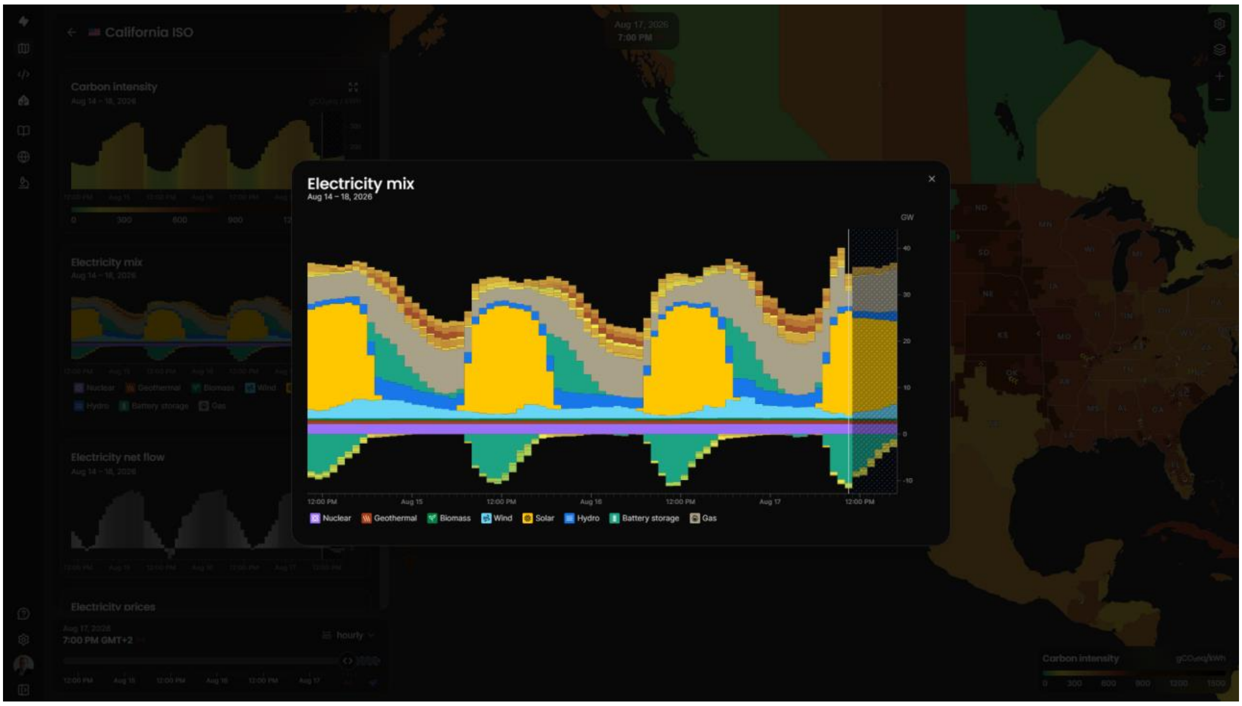

Figure 2.3: *CAISO BESS Saturation: Full Absorption Capacity Reached After Approximately 1.5–1.6 Hours, Followed by Continued Solar Generation and Curtailment.*

Fraunhofer reports that the available battery fleet reaches its full energy-absorption capacity after approximately 1.5 to 1.6 hours of charging. The CAISO profile between 14-18 August 2026 shows precisely this operational sequence: battery charging rises during the morning solar ramp, reaches its absorption limit around the solar peak and subsequently declines while solar generation remains high. The batteries were therefore effectively full before the solar-production period had ended. California's 4.5 TWh of curtailment during the first half of 2026 demonstrates the resulting system limitation.

2.6.1 The CAISO 24/7 Gas-Generation Cycle and Continuous Computing Loads

California is used here as a high-solar reference case. Large parts of the state lie within the high-irradiance US Southwest. EIA reported 25 GW of utility-scale solar capacity in CAISO in April 2026, while California accounted for approximately 34% of total US small-scale PV capacity at the end of 2025. Texas now exceeds California in utility-scale solar capacity; California nevertheless combines exceptional solar resources, one of the country's largest utility-scale PV fleets and by far its largest distributed PV base. [40]

The time-stamped CAISO observations reveal a continuous gas-fired generation layer across the daily solar cycle. At 05:00 PDT, before sunrise, gas supplied 13.4 GW, or 49.47% of available electricity. At 12:00 PDT, during the solar-production window, the gas share fell to 27.55%, but absolute gas-fired output declined by only 1.9 GW to 11.5 GW. After sunset it rose to 17.8 GW at 19:15 PDT and remained at 15.4 GW at midnight. [41]

CAISO phase	Local time	Gas output	Gas share	Installed gas capacity utilized
Before sunrise	05:00 PDT	13.4 GW	49.47%	42.21%
Solar noon	12:00 PDT	11.5 GW	27.55%	36.19%
Post-sunset evening	19:15 PDT	17.8 GW	40.21%	56.07%
Midnight	00:00 PDT	15.4 GW	49.59%	48.61%

The percentage change is larger than the physical reduction in gas-fired output: from 05:00 to noon the gas share falls by 21.92 percentage points, whereas gas generation falls by only 14.2%. Solar expands the daytime supply stack, but gas remains in operation throughout the observed cycle. For continuous computing loads, the relevant system feature is therefore not merely the midday renewable peak but the availability of dispatchable supply across every hour. The screenshots describe the regional operating environment; they do not allocate particular generators or megawatt-hours to individual data centres. [41]

2.6.2 Noon-to-Evening Comparison Across Major U.S. Data-Centre Regions

CAISO, ERCOT and PJM encompass three of the most important US data-centre markets: California's technology clusters, the rapidly expanding Texas market and Northern Virginia within PJM. These loads require continuous electricity rather than energy available only during a limited production window. Matching the observations by local time on 20 August 2026 shows substantial gas-fired generation at noon and higher output after sunset in all three systems. [41]

System	Local noon	Gas at noon	Evening local time	Gas after sunset	Increase
CAISO	12:00 PDT	11.5 GW / 27.55%	19:15 PDT	17.8 GW / 40.21%	+6.3 GW
ERCOT	12:00 CDT	30.3 GW / 35.58%	19:00 CDT	44.5 GW / 50.74%	+14.2 GW
PJM	12:00 EDT	58.7 GW / 47.29%	20:00 EDT	61.8 GW / 49.52%	+3.1 GW

The comparison distinguishes three regional structures. CAISO shows a pronounced solar-shaped daytime profile followed by a 6.3 GW evening increase in gas generation. ERCOT records the largest absolute ramp, adding 14.2 GW after the solar window. PJM is already highly gas-dependent at noon and changes comparatively little after sunset. The observations do not prove that the recorded gas generation serves data centres exclusively; they document the grid environment in which major 24/7 computing clusters operate. Full screenshots and time-zone conversions are provided in Appendix E. [41]

2.7 United Kingdom: Wind Curtailment

During the first five months of 2026, the United Kingdom curtailed over 6 Terawatt-hours (TWh) of wind generation. This volume equals the annual electricity consumption of all UK data centers combined. Grid constraint payments to wind farm operators reached £960 million (almost £1 billion) (~$1.27 billion USD) during this period, with costs passed directly to consumers through network charges. [30]

2.7.1 Germany: 24.13% of Available Offshore Output Curtailed

Germany's official first-quarter 2023 data provide a particularly clear distinction between installed capacity, generated electricity and usable energy. Bundestag Printed Paper 20/9016 reports 6,804,208 MWh of realised offshore-wind generation and a further 2,164,022 MWh of reductions ordered through grid-congestion management. The combined realised and curtailed generation potential was therefore 8,968,230 MWh. [34]

2,164,022 MWh / (6,804,208 MWh + 2,164,022 MWh) = 24.13%

Almost one in four technically available offshore-wind megawatt-hours was therefore not converted into usable electricity during the quarter. Strictly speaking, already generated electricity was not destroyed; turbines were curtailed and available wind energy was left unconverted. The economically relevant result is nevertheless a loss of generation potential. The Federal Government explicitly attributed the intervention to grid bottlenecks and stated that faster electricity-network expansion was required. [34]

The same response records 8,071 GWh of renewable-energy curtailment for 2022: approximately 5,682 GWh, or 70%, caused by transmission-grid constraints and 2,389 GWh, or 30%, caused in distribution networks. This evidence strengthens the NRG4NOW system diagnosis: where electricity cannot be transmitted or absorbed, a parallel controllable storage and logistics pathway can preserve part of the otherwise lost energy opportunity. [34]

2.8 Germany: PV, Wind and the January Paradox

Germany combines a large summer PV peak, substantial wind capacity and a pronounced winter solar deficit. The following evidence is grouped to distinguish seasonal availability from annual nameplate additions.

2.8.1 The January 2026 Winter Paradox

Electricity Maps recorded only 1.67 TWh of German solar generation in January 2026, equal to 3.36% of 49.7 TWh of electricity available during the month. The displayed monthly utilization was 2.23% of the listed production potential. The winter result demonstrates why installed summer peak capacity and short-duration batteries cannot substitute for a seasonal reserve. [8]

More grid and more BESS without seasonal storage remain a dead end: copper can transfer simultaneous power and lithium-ion batteries can shift hours, but neither can preserve a summer or autumn surplus for a January deficit.

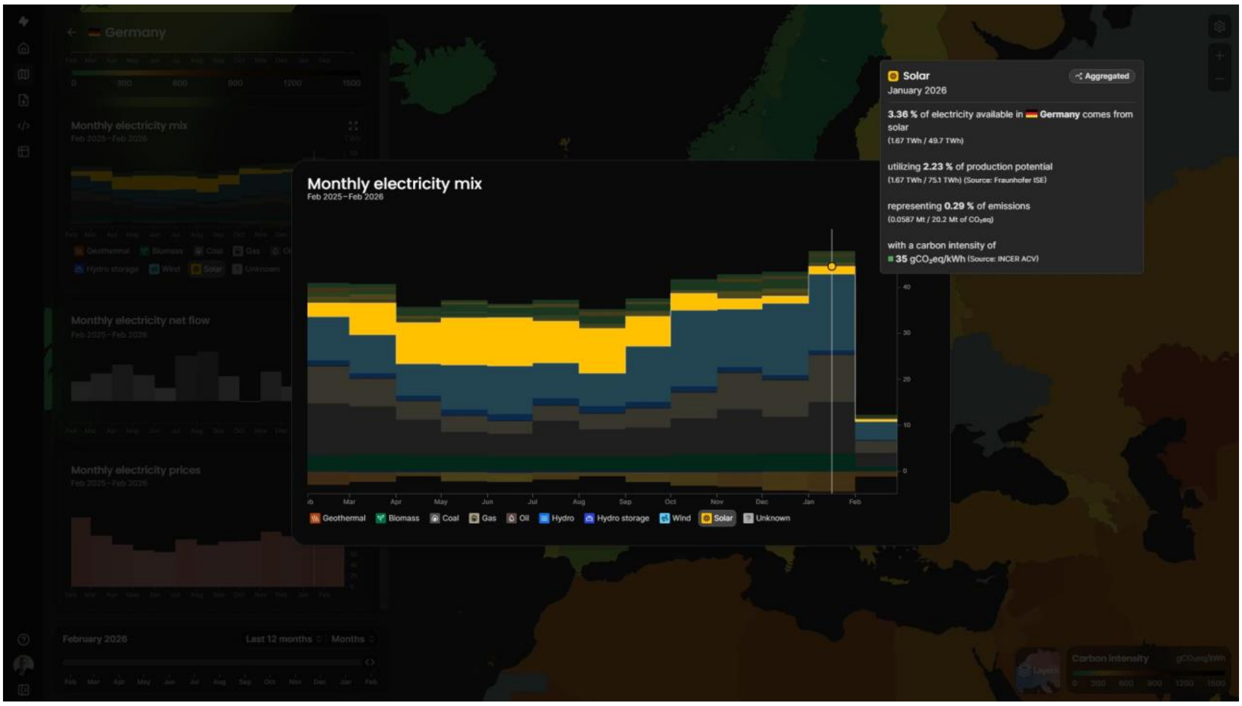

Figure 2.4: *January 2026 monthly mix: 1.67 TWh solar out of 49.7 TWh, despite a very large installed PV fleet.*

2.8.2 Wind Generation During Continued Capacity Expansion

The German case illustrates why installed capacity alone cannot establish system value. Despite continued wind-capacity additions, annual wind generation in the dataset supplied for this White Book declined from 143 TWh in 2023 to 138 TWh in 2024 and 133 TWh in 2025. [8]

Year	Wind generation	Change from prior year [8]
2023	143 TWh	baseline
2024	138 TWh	-5 TWh / -3.5%
2025	133 TWh	-5 TWh / -3.6%

The decline cannot be attributed to one mechanism without plant-level dispatch and weather-normalized data. Weather variability, maintenance, market dispatch, grid constraints and curtailment can all contribute. The synchronized August observations nevertheless justify testing the specific hypothesis that wind output is reduced during periods of high solar feed-in even when meteorological wind remains available. The separate data annex should document this relationship before it is stated as a causal annual result.

2.8.3 The Price-Spread Paradox: From Cannibalisation to Flexibility

Negative prices are a market signal that available generation, demand, grid capacity and flexibility do not coincide in time and location. Green Energy Tools' tracker, derived from ENTSO-E Day-Ahead Sequence 1 data for the DE-LU bidding zone, records 573 equivalent negative-price hours in 2025 and projected 700 to 900 hours for 2026 as of 26 June. Since October 2025, the source counts 15-minute intervals and divides their number by four; it reports that the 2026 events were concentrated mainly between 10:00 and 15:00 during solar-production hours. [35]

The widening market range is structural. After the harmonised Single Day-Ahead Coupling minimum clearing price of -500 EUR/MWh was reached on 25 April 2026, EPEX SPOT announced a lower floor of -600 EUR/MWh from 28 April. In the same market, evening scarcity can produce prices above 500 EUR/MWh. An annual average price therefore conceals the operational value of temporal flexibility. [36]

For an arbitrage asset, the economically relevant variable is the spread rather than the sign of the charging price. Co-located batteries and hybrid PPAs can monetise intraday differences; LevelTen's Q2 2026 index reported that 29% of the European PPA offers in its dataset came from Spanish and German hybrid projects. This is evidence of a changing contracting model, but it does not mean that 29% of all European PPAs contain batteries. [37]

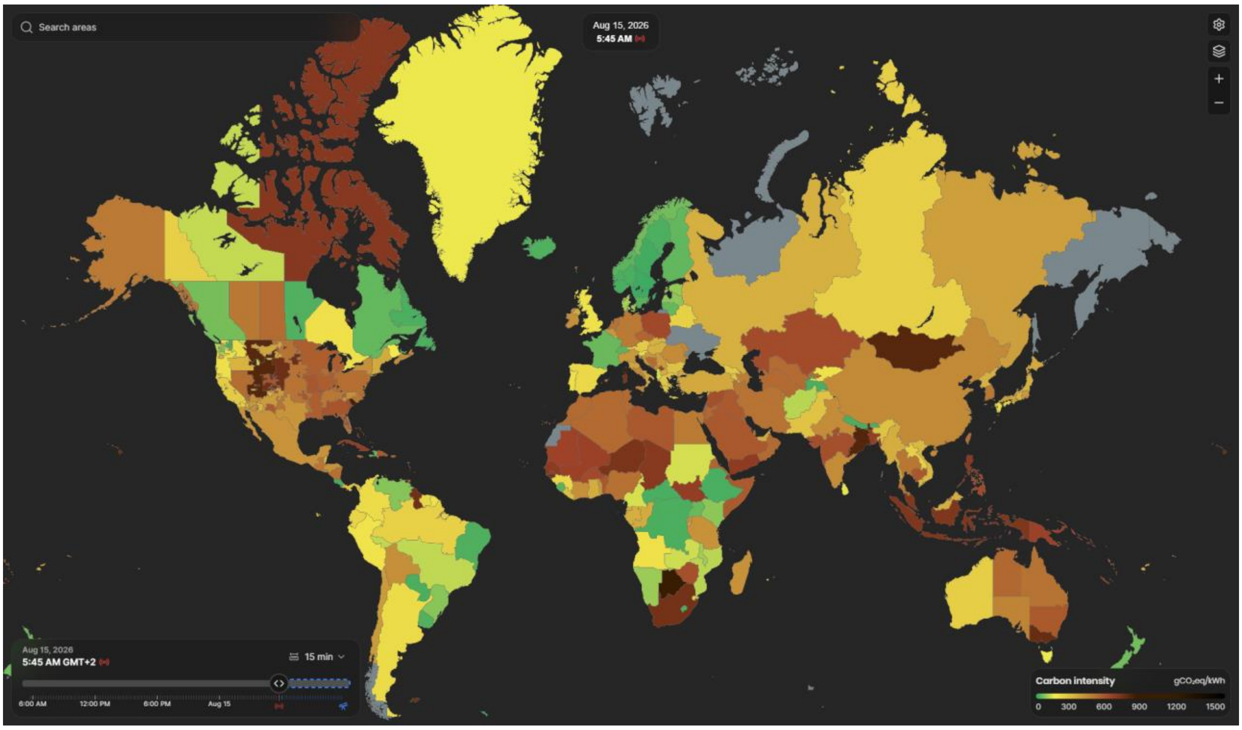

Figure 2.5: *World Map as of August 15, 2026: CO_2 Intensity and Electricity Mix Global empirical snapshot showing that high renewable nameplate capacity additions coexist with substantial fossil fuel generation and regionally divergent carbon intensity.*

The NRG4NOW extension is temporal and spatial. BESS can shift electricity from noon into the evening, while molecular storage and logistics are intended to move recoverable energy across days, seasons and constrained regions. Dan Brown's formulation that value is 'changing address' is therefore adopted here as a market interpretation, not as a substitute for measured price and curtailment data. [35–38]

2.9 China: Capacity Is Not Delivered Energy

China illustrates the distinction between installed power and delivered energy at unprecedented scale. According to the National Energy Administration, the country's installed fleet reached approximately 1,270 GW of solar capacity and 680 GW of wind capacity by 30 June 2026. [21]

Electricity Maps' synthetic 2025 annual dataset, based on Ember yearly electricity data, reports approximately 1,180 TWh of solar generation and 1,130 TWh of wind generation out of approximately 10,400 TWh of electricity available in China. Solar and wind therefore supplied approximately 22.20% of annual electricity. [22]

China evidence point	Solar	Wind [20–22]
Installed capacity, 30 June 2026	approximately 1,270 GW	approximately 680 GW
Generation during 2025	approximately 1,180 TWh	approximately 1,130 TWh
Share of 2025 electricity available	11.36%	10.84%
Capacity implied by EM production potential	approximately 888 GW	approximately 522 GW

The Electricity Maps production-potential denominators of 7,780 TWh for solar and 4,570 TWh for wind imply capacity values of approximately 888 GW and 522 GW respectively. These values correspond broadly to China's end-2024 fleet and do not represent the official end-2025 or June-2026 capacity. [20–22]

China therefore demonstrates the Gigawatt Fallacy particularly clearly: rapid additions of nameplate capacity do not translate proportionally into annual delivered energy. Transmission can move electricity geographically, but it cannot carry midday or summer production into a winter deficit. The detailed data reconciliation, system analysis and proposed pilot-corridor application are presented in Annex China.

2.9.1 Seasonal Buffering as a Complement to UHVDC

NRG4NOW does not replace China's national grid or its UHVDC corridors. It adds a temporal logistics function for qualified energy that cannot be transmitted or consumed when generated. The capacity and generation reconciliation, system-function comparison, evidence limits and proposed pilot-corridor application are developed separately in Annex China.

2.10 Industrial Consequences of Persistent Overcapacity

Renewable energy production can become decoupled from actual energy demand. This market analysis by the Financial Times describes China's production capacity for solar modules as approximately 1,000 GW per year. However, global demand for this product far exceeds production capacity. More than 40 Chinese solar manufacturers have gone bankrupt, been acquired, or been delisted from the stock exchange. One-third of the workforce at the five largest remaining companies has been laid off. [9]

Wind presents the opposite industrial risk: GWEC/BCG scenarios indicate that deployment could fall roughly 650 GW short of 2030 net-zero trajectories because manufacturing, ports, vessels, permitting and grid integration cannot expand at the required rate. Together the cases show why nameplate targets and factory output must be connected to a deliverable-energy and logistics architecture. [10]

CHAPTER 3: ENERGY LOGISTICS INSTEAD OF NEW GENERATION

3.1 The Ford Assembly-Line Principle

The preceding system diagnosis identifies a large existing energy base that is not fully usable where and when demand arises: active curtailment, pre-emptively throttled generation, renewable projects waiting for grid connection and molecular resources stranded by missing transport infrastructure. At the same time, hyperscalers respond to delayed grid access by planning dedicated new power plants. This treats the absence of an energy-delivery system as if it were primarily an absence of energy production. [4, 6, 8, 17]

The Ford principle offers the opposite perspective: build the assembly line instead of repeatedly building another energy factory. Ford's decisive contribution was not a new raw material, but a production and logistics system based on standardized interfaces, repeatable components, synchronized operations and continuous throughput. NRG4NOW transfers this principle to energy. Its 'assembly line' collects energy at the point where it would otherwise be curtailed, blocked or stranded; converts it into a qualified molecular inventory; stores and transports it in standardized modules; verifies it through the docking handshake; and converts it into electricity, heat, cooling or feedstock at the destination.

For hyperscalers, the strategic implication is direct: before commissioning another isolated gas-fired, nuclear or renewable power station, they can examine whether existing energy can be made available through a standardized logistics chain. The same infrastructure can connect renewable producers, offshore fields, ports, brownfield power-plant sites, OEMs and data centers. It can therefore address several constraints simultaneously—curtailment, interconnection queues, seasonal mismatch, stranded generation, local resistance to new power plants and the risk that accelerator hardware becomes obsolete while waiting for grid access.

NRG4NOW consequently does not begin with the question 'Which new power plant must be built?' It begins with the question 'Which existing energy is currently inaccessible, and which standardized sequence of collection, storage, transport, docking and conversion will deliver it to the load?' This is the meaning of the Ford assembly-line principle within the NRG4NOW architecture: logistics before additional production capacity.

3.2 The Physical Limit of Copper

Copper and aluminum transmission lines are point-to-point conduits requiring real-time balance. Power injected at a generation node must be consumed at a load node in the same instant.

When generation variance exceeds grid regulation limits, transmission capacity becomes a bottleneck. The NRG4NOW framework decouples electron generation from real-time consumption by converting electrical energy into pressurized, transportable molecules. Molecules can be buffered, stored, and routed independent of grid timing.

3.3 From Stationary Infrastructure to Rotating Energy Inventory

Conventional infrastructure assumes that energy remains tied to a fixed plant, line, pipeline or cavern. NRG4NOW introduces a rotating inventory: qualified modules can be filled at one node, stored above ground, moved by road, rail, barge or ship, and dispatched at another node. The asset that rotates is the certified energy inventory; generation and demand sites do not have to be joined by one permanently dedicated corridor.

3.4 Compression Work as a Controllable Energy Sink

Texas absorbs its solar peak through synchronized cooling work; Portugal absorbs part of its peak through pumped-hydro work. Northern and Central Europe require an equivalent controllable mechanical sink. NRG4NOW assigns this role to electrified compressors operating Power-to-Pressure (P2P) cycles.

5. Conversion of curtailment into density: instead of disconnecting wind and solar plants during congestion, surplus electricity drives modular compressors at ports, coastal landing points and brownfield Giga-Hubs.
6. Multi-season effect: 250-bar compressed molecules preserve the stored molecular energy without electrochemical self-discharge and remain available across the seasonal interval.
7. Winter bridge: molecules compressed during summer PV peaks and autumn wind surpluses become a strategic reserve during January nights and Dunkelflauten.

Compressors therefore become Europe's functional counterpart to the American air-conditioning load, with one decisive difference: the work is preserved as dispatchable molecular inventory rather than consumed only as immediate cooling.

3.5 The Open Modular Storage Layer Versus Salt Caverns

Conventional seasonal gas-storage logic depends on salt caverns, depleted reservoirs or other suitable geology. The storage location is dictated by subsurface conditions rather than by the point of curtailed generation or the final demand centre.

System attribute	Geological salt cavern	Qualified above-ground module fleet
Location	Restricted to suitable geology	Industrial land and brownfield hubs
Deployment	Large, site-specific project	Modular and incrementally scalable
Access	Fixed pipeline connection	Road, rail, barge and ship rotation
Demand proximity	Geology determines distance	Storage can be deployed at demand nodes
Reversibility	Immobile asset	Modules can be unbolted and rerouted
Architecture	Monolithic reservoir	Standardized cyber-physical module identity

NRG4NOW does not prescribe a proprietary vessel geometry, material or manufacturer. Its storage layer accepts certified stationary and mobile pressure modules whose carrier class, working pressure, usable inventory, cycle life, leakage performance, connector geometry, digital identity and transport approvals satisfy the published conformance profile. TSTM remains the originating reference concept for modular above-ground storage, but it is neither a mandatory product nor the sole basis of the architecture. [1, 23–26, 29]

This open approach converts industrial brownfields, ports and former power-plant sites into geographically independent energy barns. Capacity can be added in qualified increments and placed where surplus is captured or where firm energy is required. Geological storage remains valuable at suitable sites; modular above-ground storage addresses locations, pilot scales and routing requirements for which a cavern and its pipeline network are unavailable or too slow.

At 250 bar, compressed methane contains approximately 2.0–2.5 MWh of gross chemical energy per cubic metre of vessel volume, depending on temperature, gas composition, compressibility and the selected heating-value convention. The figure is not net electrical output: usable pressure swing, compressor auxiliaries and conversion efficiency must be deducted. [29]

Storage metric	Indicative volumetric value	Boundary [29]
Compressed methane at 250 bar	approximately 2.0–2.5 MWhth/m³	Gross chemical energy in vessel volume
Containerized BESS	project-specific; compare on declared system footprint	Usable electrical energy, including enclosure and auxiliaries

A fair comparison must keep thermal and electrical units separate and must state whether vessel-only, container-level or full-site volume is being compared. For the Moorburg reference, 50 modules × 13.28 GWh equals 664 GWh gross chemical inventory; any lower figure must be identified as usable or reserved inventory rather than gross arithmetic capacity. [29]

3.6 Carrier-Specific System Boundaries

The architecture standardizes interfaces; it does not make different energy carriers physically interchangeable. Every deployment must define a carrier-specific source, pressure or phase, conversion route, safety class, emissions boundary and end use.

Carrier / stream	Source and transport form	Target conversion	Principal validation requirement
CNG / methane	Fossil, biogenic or synthetic methane; high-pressure module	Engine, turbine, SOFC or pyrolysis	Origin, methane leakage, pressure safety and full carbon balance
LPG / NGL	Offshore or refinery stream; pressurized or liquid phase	Industrial feedstock or energy conversion	Composition, refrigeration, hazardous-goods logistics
Hydrogen	Electrolysis or methane pyrolysis; dedicated pressure system	Fuel cell, turbine or industry	Embrittlement, volumetric density, leakage and conversion losses

Carrier / stream	Source and transport form	Target conversion	Principal validation requirement
Synthetic methane	Renewable power plus certified carbon source	Gas conversion or destination pyrolysis	Additional conversion step, cost and carbon provenance

3.7 Two Distinct Energy-Input Pathways

The architecture must not blur electricity-derived inventory with gas that already exists as a molecular stream. They enter the system through different equipment, balances and emissions boundaries.

Pathway	Input and process	Output and accounting boundary
A — electrical surplus	Metered curtailed, throttled or otherwise qualified electricity drives compression or a separately qualified conversion process	Molecular inventory created or densified by that electricity; electricity consumption, carrier acquisition, conversion loss and emissions are recorded
B — stranded molecular stream	Associated, flared, constrained or otherwise stranded gas is conditioned, compressed and loaded	Existing molecules become transportable inventory; upstream origin, conditioning, leakage, compression power and end conversion are recorded

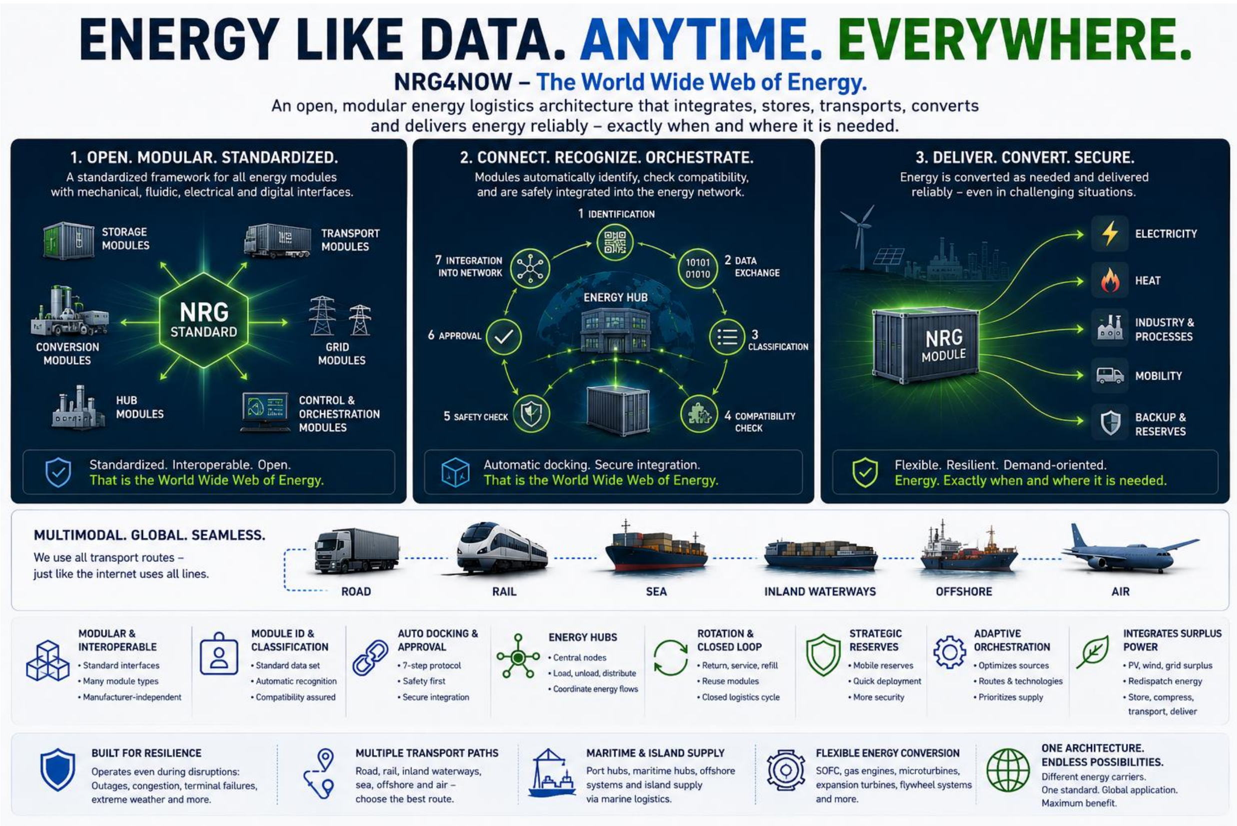

Figure 3.1: *Central World Wide Web of Energy architecture: standardized modules, machine-readable identity, docking handshake, multimodal transport, hubs and carrier-specific conversion paths.*

3.8 Historical Precedents for Maritime CNG Transport

Marine transport of compressed natural gas is not a new proposition. The 2015 World Bank Group / Tractebel report, prepared under the Global Gas Flaring Reduction programme, surveyed a substantial body of earlier development work and compared six marine CNG families: CETech, EnerSea Votrans, Sea NG Coselle, TransCanada CNG, Knutsen PNG and Trans Ocean Gas. Their containment concepts ranged from steel or composite cylinders to coiled line-pipe systems and multi-element gas containers. The report therefore establishes that maritime CNG had already progressed far beyond a speculative idea and had been examined as a route for monetizing associated and stranded gas without constructing a complete LNG chain. [13]

The report divided the CNG chain into production, transportation, receiving and storage. It found onshore CNG to be established, while marine CNG remained unproven in commercial operation as of 2015. Nevertheless, it assessed marine CNG as potentially economic for large volumes and transport distances of up to approximately 2,000 nautical miles, depending on the project. It also identified the fundamental attraction relative to LNG: lower investment in fixed liquefaction and regasification facilities. At the same time, transport cost, containment mass, loading and unloading infrastructure, gas composition, distance and project-specific utilization governed the outcome. [13]

3.8.1 Why Earlier Projects Failed

Private technical correspondence retained in the author's archive reconstructs why several technically advanced marine-CNG initiatives did not reach durable commercialization. In the mid-1990s, Enron reportedly spent more than USD 10 million developing a CNG shipping route from Oman to the Dabhol power plant. An initial bottle-ship concept was rejected because thousands of vertical cylinders required excessive spacing, connections, controls and manifold complexity. The resulting volumetric inefficiency and system cost led to the Coselle concept: a much larger pressure vessel formed from comparatively inexpensive line pipe. [15]

The Coselle system subsequently moved through Williams and Sea NG. It was tested and obtained ABS approval, and a proposed Mexican project with Enbridge reportedly passed its economic and regulatory hurdles. The project nevertheless stopped when Enbridge's board declined the risk and no suitable shipyard could be secured without financing a dedicated Coselle production facility. The later Optimum system was also tested and approved by ABS, but its corporate owner changed direction toward compressed-hydrogen shipping before a CNG project could be launched. [15]

Historical barrier	Observed effect	NRG4NOW design response [13, 15]
Thousands of separate bottles	Low space efficiency; many valves, connections and manifold controls; high vessel cost	Qualify repeatable containment modules and minimize connection classes through standardized interfaces
Specialized manufacturing	A novel containment system required a dedicated production facility that a small developer could not finance	Use an open, multi-OEM qualification framework rather than dependence on one proprietary fabrication route
Classification and weight	Temporary land-code vessels can require higher safety factors and greater mass; vessel stability and gas capacity suffer	Define the installation case early; permanently integrated systems may be classed under marine CNG rules after detailed analysis and testing
Corporate timing	Financial distress, partner withdrawal and strategic pivots terminated technically advanced programmes	Build a coalition in which ports, producers, offtakers, OEMs and capital providers share the delivery chain
Project interfaces	Import and export permits, local approvals and contracts between producer and purchaser overwhelmed small developers	Treat permits, title, custody transfer, docking, conversion and offtake as one stage-gated architecture

3.8.2 JAYANTI BARUNA: Evidence of Commercial Operation

The historical sequence did not end with failed development programmes. A Petromindo report dated 15 October 2025 documents the commercial operation of JAYANTI BARUNA, described by PLN as the world's first CNG vessel. The ship transports compressed natural gas from Gresik in East Java to Lombok in West Nusa Tenggara. At Lombok, the gas is decompressed and supplied to the 130–150 MW Lombok Peaker Gas-Fired Power Plant. This is the operational realization of the Gresik-Lombok chain that the 2015 World Bank / Tractebel study had still described as a project under construction and planned for operation in 2016. [13, 14]

JAYANTI BARUNA operating evidence	Reported result [14]
First gas-in	September 2024
Voyages completed by October 2025	25 trips
Gas delivered	424,246.96 MMBtu
Vessel CNG capacity	23.38 MMSCF
Diesel displaced	Equivalent to 7,514 kilolitres
PLN-reported annual saving	Approximately Rp 242.4 billion
Reported delivered-energy comparison	CNG carrier: USD 14–15/MMBtu; mini-LNG: USD 18–20/MMBtu; diesel: approximately USD 24/MMBtu

The Gresik CNG plant was reported at 28.62 MMSCF, the Lombok receiving plant at 22.62 MMSCF and the vessel at 23.38 MMSCF. The complete project, including the vessel, was developed by a consortium comprising Shijiazhuang Enric Gas Equipment, PT Enviromate Technology International and the Ocean Engineering Design & Research Institute of CMIC. Petromindo reported a contract value of USD 136.7 million plus Rp 244.7 billion. [14]

JAYANTI BARUNA is therefore more than an isolated vessel reference. It closes the evidentiary gap between decades of engineering studies and an operating marine-CNG logistics chain. The case does not prove that every marine-CNG route is economic, because cost depends on distance, utilization, gas quality, terminals, financing and the displaced fuel. It does prove that compressed gas can be loaded, transported by sea, received, decompressed and supplied continuously to an island power system at a reported cost below the available mini-LNG and diesel alternatives. [14]

3.8.3 Commercial Architecture Before Fleet Scale

The historical record does not show that marine CNG is physically impossible. It shows that technical approval alone was insufficient. Earlier projects concentrated large technical, manufacturing and contractual risks inside a single new vessel programme and often depended on a small developer overcoming every interface at once. A purpose-built CNG carrier can provide better transport efficiency than removable modules on a standard vessel, while removable modules may provide a lower-threshold pilot route. That trade-off must be measured rather than obscured. [13, 15]

NRG4NOW therefore begins with the two conditions identified in the correspondence as missing from earlier attempts: either a deliberately simple pilot or a coalition large enough to overcome the full project chain. The proposed architecture treats the vessel as one element of a mobile, scalable pipeline network connecting gas source, conditioning, compression, containment, marine transport, receiving hub, pipeline injection or local conversion. Its claim is not that modularity automatically makes marine CNG economic. Its claim is that standardization, multi-vendor qualification and staged allocation of risk can prevent a viable energy route from failing solely because one developer must finance a new containment factory, ship, terminal, permits and offtake agreement simultaneously. [15]

CHAPTER 4: THE OPEN REFERENCE ARCHITECTURE

4.1 NRG4NOW as the Linux of Energy Logistics

NRG4NOW operates as a vendor-neutral abstraction layer for physical energy distribution. In software engineering, the Linux kernel abstracts underlying system hardware—allowing diverse application software to execute seamlessly across heterogeneous silicon architectures. Similarly, the TCP/IP suite abstracts physical transmission media, routing data packets reliably whether the underlying link is fiber-optic, copper, satellite, or wireless.

NRG4NOW translates this abstraction paradigm directly to energy logistics. By decoupling primary generation sources (solar, wind, hydro, geothermal, nuclear, or natural gas) from the final point of consumption, the framework abstracts both energy carriers and physical transit routes. Energy is packaged into cyber-physical, standardized containers that communicate their status autonomously, enabling adaptive routing, dynamic load-balancing, and seamless multimodal handoffs across land, sea, and air.

4.2 Patent Core DE 10 2026 003 523.6

The legal, structural, and technological foundation of the Alternative Architectural Framework is anchored in patent application DE 10 2026 003 523.6, filed with the German Patent and Trade Mark Office (DPMA) on June 22, 2026 ("Standardisierte Energie-Logistikarchitektur für molekulare Energieträger mit austauschbaren Speicher-, Transport-, Umwandlungs-, Netz- und Hub-Modulen"). [1]

The patent specification does not claim an individual containment vessel or a singular turbine design. Instead, it establishes an open, interoperable system architecture defined by five core pillars: [1]

8. Universal Interface Standardization: Definition of unified mechanical, fluidic, electrical, and digital data interfaces across all system components. This encompasses storage modules, transport modules, energy conversion modules, grid interface modules, energy hubs, and orchestration units. [1]
9. Machine-Readable Module Identity & Classification: Every module carries a cryptographically verifiable, machine-readable dataset. This profile encodes unique module identity, module type, energy carrier classification, operating pressure class, nominal power class, safety/hazard classification, communication protocol, active operational state, and maintenance history. [1]
10. Closed-Loop Rotational Logistics: Replacing static, location-bound refueling with an active rotational logistics loop. Depleted modules are not refilled on-site under variable field conditions; instead, they are uncoupled, rotated back to dedicated Energy Hubs for automated inspection, servicing, and recharging, while pre-filled modules rotate into service. [1]
11. Multimodal Routing & Infrastructure Resiliency: Modules are dimensioned and certified for seamless transit across road, rail, inland waterways, maritime shipping, offshore systems, and aerial transport. If a primary transit corridor or terminal is obstructed, the orchestration layer automatically re-routes modules via alternative modes. [1]
12. Technology-Agnostic Conversion: The stored molecular energy can be dispatched at the point of consumption through interchangeable conversion modules—including Solid Oxide Fuel Cells (SOFC), internal combustion gas engines, microturbines, expansion turbines (capturing cold energy), or flywheel buffers. [1]

4.3 The Autonomous Seven-Step Handshake

When a module physically docks at an Energy Hub or a consumption facility, the control architecture executes an automated 7-step safety and integration handshake prior to opening fluidic or electrical valves: [1]

THE AUTONOMOUS 7-STEP MODULE DOCKING & INTEGRATION HANDSHAKE

This protocol eliminates human operator error during high-pressure fluid transfers and ensures that incompatible energy carriers or pressure classes can never be cross-connected. [1]

Step	Protocol state	Function
1	Physical identification	Read module identity and carrier class
2	Data exchange	Transfer pressure, temperature, state and provenance
3	Classification	Assign pressure, safety and operating class
4	Compatibility	Verify mechanical, fluidic, electrical and digital interfaces
5	Safety check	Validate seals, valves, isolation and emergency state
6	Release	Authorize coupling and controlled transfer
7	Grid / hub integration	Register module in orchestration and dispatch

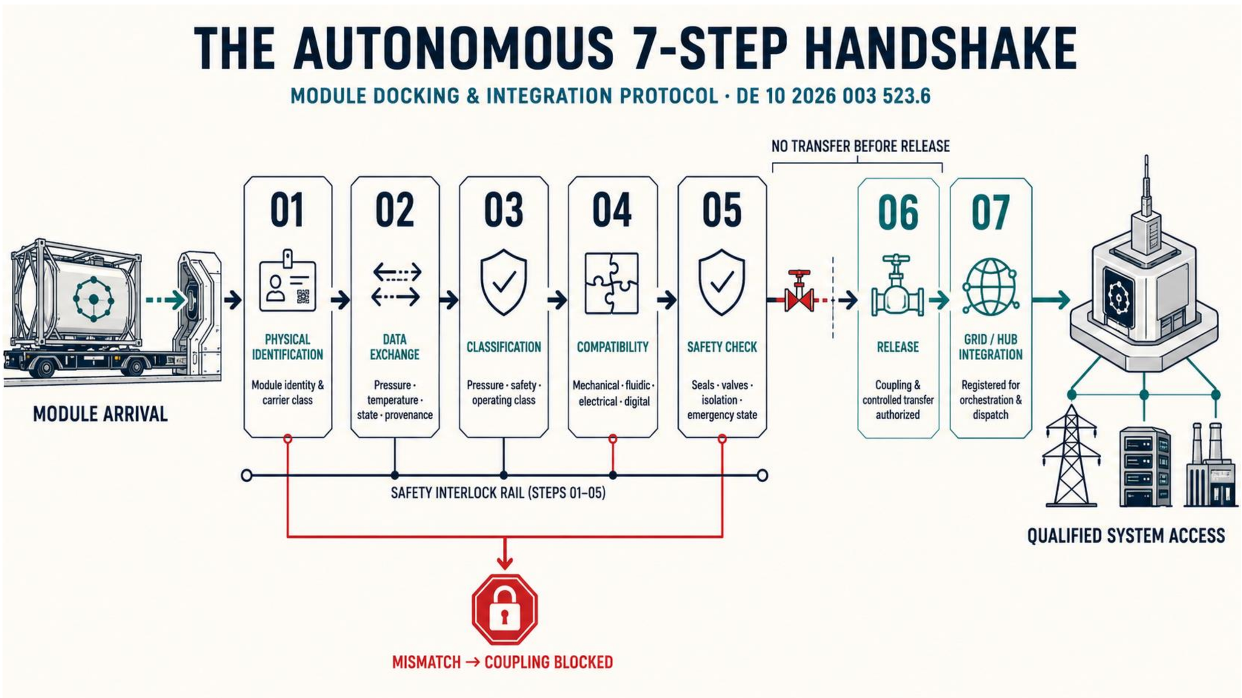

Figure 4.1: *The Autonomous 7-Step Handshake (Patent Application DE 10 2026 003 523.6) – Standardized docking, mutual authentication, compatibility verification, and deterministic safety interlock rail (Steps 01–05) governing fluidic release (Step 06) and dispatch integration (Step 07) into qualified compute and grid infrastructure.*

4.4 Vendor-Neutral Module Qualification

NRG4NOW defines a conformance envelope rather than a proprietary storage product. A supplier may participate when its certified module and associated equipment satisfy the carrier-specific mechanical, safety, transport and digital-interface requirements. The commercial architecture therefore remains open to established manufacturers and new entrants. [1]

Qualification field	Minimum project record
Carrier and material compatibility	Declared gas composition, liner and structural materials, permeation and embrittlement limits
Pressure and inventory	Maximum allowable working pressure, permitted swing, residual inventory and verified usable mass/energy
Life and integrity	Cycle limit, inspection interval, fatigue basis, leakage threshold and end-of-life rule
Interfaces	Mechanical connector, valves, isolation, telemetry, digital identity and seven-step handshake profile
Mobility and approvals	Stationary, road, rail, inland-waterway or maritime class; applicable pressure and dangerous-goods approvals
Performance claim	Gross chemical inventory and net deliverable electricity stated separately for the selected conversion route

4.4.1 Illustrative Industrial Participant Landscape

The following companies illustrate an existing industrial supply base that could compete or cooperate within an NRG4NOW pilot. Their inclusion records technical market capability; it does not imply a partnership, endorsement, commitment or conformance certification by NRG4NOW. [23–28]

Capability	Illustrative manufacturers	Relevant published capability [23–28]
CNG/RNG storage and transport	Hexagon Agility; CIMC Enric	Mobile Pipeline composite modules and certified CNG tube skids/MEGC products
Hydrogen storage and MEGC	NPROXX; Luxfer Gas Cylinders	Type 3/4 vessels and transport systems at carrier-specific high pressures
Reciprocating compression	Ariel; Burckhardt Compression; HOERBIGER	Natural-gas, RNG and hydrogen compression packages and components
Maritime/LNG compression	Burckhardt Compression; Everllence	High-pressure compressor systems for LNG-carrier applications
Gas engines	INNIO Jenbacher	Natural-gas and hydrogen-capable distributed generation
Gas turbines	Siemens Energy	Industrial turbines with fuel-flexibility and hydrogen demonstration pathways
Solid-oxide fuel cells	Bloom Energy	Modular SOFC generation using natural gas, hydrogen or qualified blends

The earlier Moorburg study—50 reference modules at 13.28 GWh gross compressed-methane energy per module, or 664 GWh gross inventory—remains a historical sizing scenario only. It does not define the required module design, supplier or net electrical output. Every implementation must be recalculated from certified OEM data and a complete source-to-load balance. [29]

4.5 Compression, Loading and Transfer Interfaces

NRG4NOW separates the physical transfer interface from the storage-vessel design. Every loading or unloading operation begins with a declared carrier, pressure class, temperature envelope, allowable flow range and emergency-isolation state. A module is released only after mechanical, fluidic, electrical and digital compatibility have been verified. [1]

Interface layer	Minimum controlled parameter	Acceptance evidence
Mechanical	Connector geometry, restraint and load case	Fit, locking and structural test
Fluidic	Carrier, pressure, temperature and flow direction	Materials record, valve state and leak test
Safety	Isolation, venting, emergency shutdown and exclusion zone	HAZID/HAZOP and functional test
Custody transfer	Mass/energy quantity, composition and ownership	Calibrated metering and signed ledger event

4.6 Digital Identity, Telemetry and Source-to-Load Ledger

Each qualified module carries a machine-readable identity linked to its manufacturer, carrier class, pressure history, inspection status, usable inventory and permitted routes. Telemetry records the state transition from filling through transport, docking and conversion. The source-to-load ledger keeps electricity input, molecular inventory, auxiliaries, leakage, water, carbon origin and useful output in separate fields so that no energy or emissions claim is counted twice. [1]

4.7 Local Deterministic Safety and Fleet Orchestration

Safety decisions remain local and deterministic. Loss of connectivity, a conflicting carrier code, an exceeded pressure/temperature limit or an invalid inspection state blocks coupling without waiting for a fleet optimizer. Higher-level orchestration selects routes and inventory positions, but it cannot override the certified local interlock. [1]

Fleet optimization may use forecasts and bounded heuristics; the safety case may not depend on solving a global routing problem in real time. Reserve modules, precomputed routes and human escalation rules preserve operation when communications or optimization services fail. [1]

CHAPTER 5: NRG-ALADDIN RISK MANAGEMENT AND MODEL-BASED VALIDATION

5.1 Method and Evidentiary Status

Two independent generative-AI model series—Microsoft Copilot and Meta AI—were each given four adversarial prompts. The models were instructed to examine NRG4NOW from four roles: a hyperscaler energy strategist, an investor seeking control of the value chain, a hostile engineering reviewer and an IEC/ISO standardization authority. The purpose was not to obtain endorsement. It was to expose strategic resistance, commercial capture incentives, technical failure modes and institutional barriers before a pilot project is launched. [18, 19]

The resulting texts were deliberately written in the form of confidential internal memoranda and adopted the voices of Microsoft, Azure, Meta and standards officials. They are simulated AI outputs. They are not genuine corporate documents, do not report real conversations with named engineers, and do not demonstrate actual decisions or intentions by Microsoft, Meta, IEC, ISO or any other organization. Numerical claims inside the scenarios are retained below as stress-test inputs and are identified as assumptions requiring independent validation. [18, 19]

Prompt	Assigned perspective	Question tested [18, 19]
1	Hyperscaler strategy	Could NRG4NOW undermine monolithic power assets, exclusive grid access and proprietary energy stacks?
2	Investor / value capture	Where could value be created, captured or monopolized if curtailed and stranded energy becomes logistically usable?
3	Hostile engineering review	Which thermodynamic, material, fleet-rotation and orchestration mechanisms could cause technical failure?
4	IEC/ISO governance	Which interfaces are standardizable, and which certification, jurisdiction and institutional barriers could delay adoption?

5.2 Hyperscaler Strategic-Threat Scenario

Both models interpreted NRG4NOW not as a single product but as an architecture standard. Their common abstraction was that modular energy carriers, standardized mechanical and digital interfaces, an autonomous docking handshake, an orchestration layer and rotational logistics would 'virtualize' energy in the same way that cloud architectures abstracted compute resources. Copilot compared the concept with Kubernetes; Meta described the handshake as 'TCP/IP for energy' and the interfaces as 'USB for energy.' [18, 19]

The simulated hyperscaler position contrasted monolithic infrastructure—multi-gigawatt power plants, exclusive grid connections, proprietary energy stacks and long-duration power-purchase agreements—with large populations of interoperable modules. In the scenario, current market power rests on three barriers: access to very large pools of capital, control of scarce grid connections and the ability to manage complex proprietary projects. An open modular standard could lower all three barriers by allowing multiple OEMs, ports, infrastructure investors and regional operators to participate in one certified system. [18, 19]

Monolithic hyperscaler model	NRG4NOW model tested	Strategic risk identified by the models [18, 19]
Dedicated 1–5 GW power projects	Large fleets of 1–5 MW modules	Modular deployment could shorten time-to-power and reduce dependence on one project
Exclusive grid and pipeline access	Road, rail, inland-waterway and maritime routing	The value of a privileged connection may decline when energy can follow another route

Monolithic hyperscaler model	NRG4NOW model tested	Strategic risk identified by the models [18, 19]
Proprietary energy stack	Open interfaces and multi-vendor modules	Integration control shifts from asset ownership toward protocol governance
Complexity as an entry barrier	Automated identity, compatibility and release protocol	Ports and industrial sites can become energy hubs without reproducing a hyperscaler's entire organization
Long-term technology lock-in	Exchangeable conversion modules	New conversion technology can enter without replacing the complete logistics chain

Both model narratives offered the same three strategic responses: block the architecture, ignore it or participate in and influence the standard. Blocking was portrayed as creating delay while competitors advance; ignoring it was associated with stranded monolithic assets; participation was presented as the least damaging response because it preserves a role in protocol governance, certification and orchestration. These are model-generated strategic alternatives, not evidence that a hyperscaler has adopted any of them. [18, 19]

The scenarios also introduced an internal-adoption hypothesis. Engineers facing grid congestion, permitting delays and inflexible legacy assets could prefer modular systems because they solve immediate operational problems and resemble familiar software architecture. The texts described a possible 'shadow-IT' pathway through temporary containerized generation and storage. No evidence was supplied that such NRG4NOW experiments or conversations were actually occurring; the point of the scenario is that operational utility can drive bottom-up adoption before corporate strategy formally changes. [18, 19]

5.3 Investor and Value-Capture Scenario

The investor prompts identified three potential value conversions: curtailed electricity into molecular inventory, a grid bottleneck into a logistics service, and site dependence into modular supply. The recurring beneficiary groups were ports and terminals that become energy hubs, infrastructure funds that finance repeatable modules and hubs instead of only long linear networks, and OEMs that manufacture standardized modules rather than bespoke gigawatt-scale plants. [18, 19]

Value stream tested	Copilot scenario	Meta scenario	Required validation [18, 19]
German curtailment base	19 TWh × EUR 60/MWh = EUR 1.14 billion	Same 19 TWh and EUR 60/MWh basis	Confirm what the 19 TWh dataset includes and avoid equating every redispatch action with discarded renewable energy
Recoverable share	70% = approximately EUR 800 million/year	80% = approximately EUR 912 million/year	Pilot-derived capture rate, compression availability, energy losses, market value and operating cost
Global value	More than EUR 50 billion/year	More than EUR 50 billion/year	Country-level datasets, common definitions and a complete net-value calculation
Infrastructure economics	Hub ROI stated as 24–36 months	Hub ROI stated as about 24 months and module cost below USD 10/kWh	Vendor quotations, utilization, financing, permitting, maintenance and replacement assumptions

The divergence between 70% and 80% recovery is analytically useful: neither value is a measured NRG4NOW performance result. It defines a sensitivity range to be tested. The same applies to the claimed return periods, module costs, margins and global market value. A bankable model must begin with metered energy available at a specific node, subtract conversion and compression demand, transport and rotation costs, maintenance, residual inventory and reconversion losses, and then value the energy at the destination and time of delivery. [18, 19]

The Meta investor role additionally simulated hostile value capture: financing compatible OEMs, forcing a nominally open standard into FRAND governance, marginalizing the inventor within a consortium and using compliance requirements to concentrate market control. The language was intentionally adversarial. Its relevance is defensive: the patent owner and the proposed NRG4NOW governance framework require transparent licensing, balanced voting rights, conflict-of-interest rules, published conformance criteria and protection against capture by any single hyperscaler, investor or manufacturer. [19]

5.4 Technical Red-Team Scenario

The strongest agreement between Copilot and Meta occurred in the technical attack. Both models identified three failure domains: thermodynamic and material integrity during coupling, availability losses during carrier or role changes, and the computational complexity of orchestrating a large heterogeneous module fleet. [18, 19]

Attack surface	Copilot assessment	Meta red-team claim	NRG4NOW design response required [18, 19]
Pressure and temperature transients	Joule–Thomson effects, fatigue, seal ageing and condensation are solvable with robust sensing and active control	A mismatched high-pressure hydrogen / methane coupling could create a temperature drop of about 40 K and severe seal or leakage risk	Carrier- and pressure-class segregation; no uncontrolled cross-connection; staged pressure equalization; validated materials, valves, seals and thermal envelope
Carrier / role change	Purging, inerting and pressure testing may impose 6–12 hours of downtime	Assumed 8–12 hours and up to 30% of a 1,000-module fleet unavailable during switching	Prefer dedicated carrier fleets; define certified cleaning state machine; size reserve inventory from measured turnaround time
Fleet orchestration	Multi-depot, multi-carrier, real-time and safety-critical routing is NP-hard but approachable with heuristics, edge AI and prioritization	A 10,000-module vehicle-routing problem cannot be fully optimized within a four-minute grid emergency	Separate safety control from economic optimization; use local deterministic fallback, precomputed routes, reserve modules and bounded heuristic dispatch

The numerical attack values—40 K temperature change, leakage above 5%, 30% fleet unavailability, EUR 500,000 per day of idle capital and a four-minute blackout window—were asserted inside the Meta role-play and were not supported by calculations or test data in the supplied PDF. They must therefore be treated as test cases, not findings. Their engineering value lies in defining what a demonstrator must measure: transient temperature and pressure, seal leakage by carrier and cycle count, material fatigue, purge duration, module turnaround, reserve requirement, routing latency and fail-safe performance. [19]

The two models nevertheless support a concrete architecture decision. The seven-step handshake cannot itself compensate for an invalid physical connection. Compatibility rules must prevent incompatible carrier, pressure, material and connector classes from reaching the release state. Likewise, the orchestration platform must never be the sole blackout-protection layer. Local governors, isolated emergency logic and immediately available reserve modules must maintain the load while route optimization continues at a slower planning horizon. [18, 19]

5.5 Standards and Governance Scenario

Both standardization prompts concluded that important NRG4NOW elements are capable of standardization. Copilot listed mechanical interfaces, pressure and safety classes, the communications protocol, the identification schema and the carrier-role-change state machine. Meta mapped possible starting points to existing families such as ISO 15500 for vehicle gas-system components, OPC UA or CAN-based data exchange, IEC 60079 for explosive atmospheres, IEC 62282 for fuel-cell systems, IECEx conformity assessment and classification-society rules for maritime deployment. [18, 19]

Standardization object	Preliminary route identified in the scenarios	Principal obstacle [18, 19]
Mechanical and fluidic coupling	Extend or profile existing pressure-system connector and vessel requirements	Different carriers, pressures, temperatures and jurisdictions
Digital identity and handshake	CAN / OPC UA-compatible profile with machine-readable module identity	Cybersecurity, semantic interoperability and liability for an incorrect release
Safety classification	Map carrier and operating classes to explosive-atmosphere and pressure-equipment rules	Overlapping technical committees and national implementation
Conversion modules	Reference applicable fuel-cell, engine, turbine and electrical interconnection standards	No single standard covers the complete modular chain
Conformance system	Cross-vendor reference implementation, interoperability tests and accredited certification	Limited laboratories, cost, waiting time and governance legitimacy

The institutional red-team scenario then examined how standardization could be used to delay or capture the architecture: disputes over technical-committee jurisdiction, rejection of a new-work proposal, selective references to existing standards, slow conformity assessment and denial of insurance or port access to non-certified modules. The scenario used provocative language and invented instructions; it is not evidence of misconduct by IEC or ISO. It does, however, identify a genuine deployment dependency: without an agreed conformity route, insurers, financiers, ports and authorities may refuse an otherwise functional module. [18, 19]

The appropriate response is a dual-track standardization strategy. NRG4NOW should develop a technically complete reference implementation and pilot evidence while simultaneously mapping every interface to existing statutory and consensus standards. A transparent pre-standard or publicly available specification can support early interoperability, but it cannot replace pressure-equipment approval, hazardous-goods law, maritime class, environmental permits or national operating authorization. [18, 19]

5.6 Cross-Model Synthesis

Cross-model finding	Architectural implication	Required evidence before scale-up [18, 19]
Open modular logistics could weaken monolithic infrastructure advantages	Keep interfaces vendor-neutral and governance multi-stakeholder	Cross-vendor docking and substitution demonstration
Curtailed and stranded energy may form a new value pool	Meter the complete source-to-load chain	Node-specific energy balance, cost, emissions and revenue data
Thermodynamic compatibility is the primary safety gate	Encode carrier, pressure, material and thermal classes before release	Cycle, leakage, transient and failure-mode testing
Role changes can destroy fleet availability	Use dedicated fleets where practical and model reserve inventory	Measured purge, inspection and turnaround times
Global optimization cannot carry the safety case	Separate deterministic local protection from fleet-level optimization	Failover demonstrations under communication and routing failures
Standardization can enable interoperability or concentrate control	Publish conformance rules and balance alliance voting rights	Independent certification route and governance charter

The combined outcome is neither an endorsement nor a technical rejection. The model exercise converts broad claims into a pilot agenda. It shows why NRG4NOW may be strategically attractive to hyperscalers while simultaneously threatening their sunk assets and platform control; why investors may finance the architecture while attempting to capture it; why engineers may value its modularity while demanding hard thermodynamic evidence; and why standards bodies are indispensable even though a slow process can outlast the commercial window. [18, 19]

For the White Book, the decisive result is operational: every benefit asserted by the strategic and investor scenarios must be paired with a measurable engineering and governance test. Time-to-energy must be demonstrated, not promised. Recoverable curtailment must be metered, not inferred from nameplate capacity. Interoperability must be proven across manufacturers. Safety must remain valid through repeated pressure cycles and failure states. Alliance governance must prevent the open standard from becoming the proprietary control layer of a single winner. [18, 19]

5.7 Four Predictive Risk Domains

NRG-ALADDIN is the proposed orchestration and predictive-risk layer for the module fleet. It evaluates four interconnected domains before deployment and continuously during operation:

Risk domain	Observed variables	Operational response
Thermodynamic integrity	Pressure gradients, cycle count, temperature, fatigue and seal condition	Block release, route to inspection, isolate or derate module
Meteorological and maritime	Wind, wave, visibility, seasonal patterns and port accessibility	Adjust floating-buffer inventory and shuttle schedule
Multimodal logistics	Port, rail, road and inland-waterway capacity or outage	Re-route modules and rebalance hub inventory
Socio-political site risk	Zoning, water stress, acoustic limits, public filings and local sentiment	Compare sites before commitment and trigger engagement or relocation

The system is a design framework, not a deployed autonomous decision engine. Its performance must be validated against real fleet data, documented failure modes, human authorization rules and the legal responsibilities of operators and regulators.

5.8 Socio-Political Site Risk

The 2026 political constraint has expanded beyond isolated project disputes. More than 500 US jurisdictions have imposed bans, moratoriums or statutory blocks on large data centres. NRG-ALADDIN therefore treats water stress, acoustic impact, household-price exposure, zoning and local sentiment as primary engineering inputs before capital deployment. [6]

The platform identifies jurisdictions where grid-tethered development is likely to fail and redirects projects toward net-decoupled brownfield or port hubs before sunk capital accumulates.

5.9 The Fire Department Principle

The Fire Department Principle converts NRG-ALADDIN from a passive risk register into a bounded intervention process. When a monitored variable crosses an approved threshold, the system does not improvise a commercial or safety decision. It assembles the responsible technical, regulatory, logistics and community actors around a predefined incident state and documented authority chain.

Trigger	Immediate action	Decision owner
Water or acoustic limit at risk	Hold site commitment; verify design envelope and mitigation	Project owner and competent authority
Carrier, pressure or material mismatch	Block docking and isolate the module	Local safety controller / operator
Route, port or fleet interruption	Activate reserve inventory or approved alternate route	Logistics control and offtaker
Community or permitting escalation	Pause irreversible capital and open documented engagement	Project governance and local authority

The principle does not promise conflict-free deployment. It prevents a known engineering, environmental or social conflict from being ignored until it becomes a stranded-capital event.

CHAPTER 6: DEPLOYMENT SCENARIOS

6.1 The Four-Dimensional System Matrix

Dimension	Onshore Giga-Hub	Offshore greenfield	Offshore brownfield	Offshore capacity lock [2, 29]
Reference	Moorburg / Hamburg	Namibia / Kudu	SNS / Bacton	Norway / Sleipner
Constraint	Grid, water, NIMBY	No multibillion-dollar pipeline	Pressure mismatch / terminal removal	Pipeline saturation
NRG4NOW logic	P2P seasonal buffer	Reversible shuttle pilot	On-site conversion	Floating compute / direct export
Mobility	Road / rail / barge	Maritime rotation	Short-sea shuttle	Floating modular layer

6.2 Offshore Compute and Brownfield Compute

Parameter	Route A: Offshore compute - Bacton / SNS	Route B: Brownfield compute - Moorburg [2, 29]
Primary energy	Associated or stranded gas close to the source; offshore wind can serve compression auxiliaries	Rotating P2P molecular inventory with carrier-specific storage balance
Cooling	Seawater heat sink through permitted closed-loop or indirect systems	Existing industrial cooling infrastructure, adapted to rack and process requirements
Siting	Lower residential exposure; maritime and environmental permits still required	Existing industrial zoning and safety distances; brownfield approvals still required
Data / energy flow	Energy conversion near source; data returned by subsea fibre	Compute beside hub; grid-independent or grid-limited operating modes
Time-to-energy	12–18 month model target for modular floating deployment	18–20 month model target for brownfield conversion

The dual-route model addresses different bottlenecks. Offshore compute places a steady computational load next to gas that cannot justify a new pipeline or cannot enter a pressure-constrained terminal. Brownfield compute places the load next to established industrial land, cooling systems and transport corridors while molecular modules provide a controllable energy inventory. Neither route is permit-free; both are designed to shorten the dependency chain that begins with a new greenfield site and ends in a congested onshore grid queue. [2, 29]

A further acquisition pathway uses Subsea Access Nodes connected to existing, underutilized or stranded pipeline systems. Gas can be withdrawn through a qualified subsea interface, conditioned and compressed on a floating production or loading unit, transferred to modular carriers and delivered to a brownfield power or compute site. Where available, offshore wind may supply part of the compression load. [29]

The pathway was described in Addressing Global Energy Transition Challenges with Modular TSTM Solutions 2.1 (Zenodo DOI: 10.5281/zenodo.15586051). It is a concept-development reference, not evidence that an unrestricted connection can be made to any pipeline. Every corridor remains subject to operator consent, integrity assessment, pressure and composition compatibility, metering, class approval and permitting. [29]

6.3 Bacton and Sleipner

In the Southern North Sea, more than 30 fields feed the Bacton terminal system. Ageing satellites can become uneconomic when declining reservoir pressure, compression requirements, tariffs and the remaining life of shared infrastructure no longer support conventional pipeline delivery. The NRG4NOW case places modular compression, storage or conversion capacity upstream of that bottleneck, allowing a compute load or shuttle chain to use gas that would otherwise remain stranded or be written down. [2]

The Norwegian Continental Shelf presents a different constraint. The supplied system analysis describes Gassled throughput of 117.6 bcm in 2024 against a design scale of approximately 120 bcm. Smaller discoveries can therefore face long allocation delays at hubs such as Sleipner or Kollsnes. A qualified modular shuttle and floating buffer would not replace the pipeline system; it would create a parallel pilot route for volumes that cannot obtain a timely pipeline slot. [2]

The investor proposition is conditional but testable: gas producers monetize otherwise uneconomic molecules, offshore wind operators gain a controllable compression load, and hyperscalers obtain a firm-energy pathway with lower dependence on onshore interconnection. Data output returns through subsea fibre while the heat-rejection system uses the maritime heat sink under site-specific environmental limits. [2, 29]

6.4 West Africa and the Niger Delta

West Africa combines associated-gas flaring, constrained gathering systems, power deficits and coastal transport routes. The NRG4NOW use case is not a regional TWh extrapolation. It is a corridor-by-corridor test in which a defined molecular stream is conditioned, compressed, loaded, transported and converted at a contracted industrial or compute load. [16]

A Niger Delta pilot must document gas title, composition, current disposition, compression auxiliaries, leakage, security, port interfaces and the displaced fuel. Only the metered source-to-load result can establish whether modular logistics is preferable to reinjection, pipeline expansion, LNG, local generation or continued non-use. [16]

6.5 The Orange Basin

In January 2025, Shell recorded a $400 million impairment write-down on its Namibian deepwater acreage (PEL 39) in the Orange Basin. This impairment was not driven by geological failure, but by the complete absence of a midscale logistics architecture to monetize associated gas and NGL streams. [2]

Similarly, at TotalEnergies' nearby Venus field (3,000m water depth), the approved development plan reinjects 500 MMscfd of associated gas. At European gas reference prices ($10–$12/MMBtu), this reinjected gas represents $1.8 to $2.2 billion in unactivated annual gross value ($36 to $44 billion over a 20-year field life). [2]

The NRG4NOW qualified modular extraction layer addresses this gap at first production. Operating under low-pressure regimes (8–18 bar), modular containment units on chartered offshore vessels capture LPG and NGL fractions directly from the FPSO, buffering the fluid and shuttling it to regional coastal markets. [2]

6.6 Moorburg Reference Configuration

The Moorburg configuration is retained as a historical sizing scenario, not a mandatory product definition: 50 reference modules × 13.28 GWh equals 664 GWh of gross chemical methane inventory at 250 bar. Net deliverable electricity is lower and depends on usable pressure swing, gas composition, auxiliaries, conversion efficiency, reserve policy and availability. [29]

Revision 16 remains vendor-neutral. A brownfield hub may use any qualified storage and compression equipment meeting the common conformance profile; no TSTM geometry or single manufacturer is prescribed. [1, 23–28]

6.7 Thermodynamic Integration, Water and Acoustic Design

Pressure reduction can provide useful process cold where the carrier, inlet temperature, pressure range and Joule-Thomson coefficient are suitable. The first-order relation is:

$$\Delta T = \mu_{JT} \cdot \Delta P$$

This relation does not by itself establish the cooling capacity of a data centre. A project calculation must include mass flow, real-gas behaviour, expander or valve selection, heat-exchanger approach temperatures, rack heat density, pumping power, backup rejection capacity and the full annual operating profile. The data-centre waste heat must still be transferred to an environmental sink.

The engineering design target is no running evaporative water consumption at the rack cooling interface under the specified operating envelope. Closed-loop liquid cooling, recovered process cold and permitted seawater or dry-cooler heat rejection can reduce freshwater demand, but the site water balance must quantify make-up, treatment, cleaning and emergency requirements.

The acoustic target is compliance with the applicable industrial and residential immission limits at the site boundary. Pumps, compressors, expanders, generators, transformers and heat rejection equipment remain sound sources. Enclosures, vibration isolation, equipment selection and distance are therefore design measures to be verified by an acoustic model and commissioning measurements; 0 dB is not a physical system claim.

6.8 Pilot-Corridor and Stage-Gate Logic

Deployment begins with a pilot corridor and the G0–G6 stage gates defined in Appendix D. The sequence freezes the use case and ownership boundary before equipment selection, then closes technology, safety, permits, finance, integration and acceptance in order. A failed gate stops or redesigns the project instead of allowing an unverified national-scale claim to advance.

CHAPTER 7: THE BUSINESS CASE: TIME-TO-ENERGY

7.1 96 Months Versus 20 Months

Development phase	Legacy grid pathway	NRG4NOW pathway [4, 6, 33]
Zoning and public protest	24 months	6 months on brownfield
Interconnection queue	60 months	Not required for closed island mode; permits remain
Manufacturing and build	12 months	12 months in parallel
Integration and testing	12 months	2 months; handshake
Total Time-to-Energy	96-month reference case	20-month design target

In the era of Artificial General Intelligence (AGI) and exponentially scaling frontier models, the time required to deliver the first Megawatt-hour of electrical power—Time-to-Energy (TTE)—has emerged as the primary economic metric for capital allocation.

The NRG4NOW framework is designed to reduce dependence on grid interconnection queues and to shift, rather than eliminate, the permitting burden toward a defined brownfield, energy, transport, safety and compute approval pathway.

For the defined 100 MW brownfield reference pilot, the NRG4NOW programme targets energization within 20 months, subject to successful completion of every stage gate and regulatory approval. The commercial relevance of minimizing Time-to-Energy follows from the short innovation cycles of accelerator hardware, the high fixed cost of data-center capacity and publicly reported hyperscaler investment plans. [6, 33]

The 20-month value is a stage-gated program target for a defined brownfield pilot, not a universal construction promise. Failure to close a safety, permitting, fuel, equipment or commercial gate stops or resets the schedule.

7.2 The 100 MW Brownfield Reference Case

A 100 MW net continuous electrical load requires 2.4 GWh per day and 876 GWh per year before any distinction between IT load and total facility load. If 100 MW denotes IT load, facility demand must additionally include power-usage effectiveness. The reference case therefore sizes the complete chain from a declared net-load boundary rather than from gross module inventory.

7.2.1 From Electricity-Price Spread to Logistics Value

A large positive difference between the absorption price and the later delivery price creates a gross flexibility opportunity, not a guaranteed profit. The recoverable value must be reduced by conversion losses, storage cost, transport cost, cycling limitations, grid charges and the duration and volume of both market windows. The same accounting principle applies to batteries and to molecular inventory, but the intended storage horizon differs. [35–38]

$$Vspread = Erecoverable \times (Pdelivery - Pabsorption) - Cconversion - Cstorage - Ctransport$$

This expression is a screening equation. A project-specific model must separately specify energy losses, capacity limits, contract structure, operating hours, pressure cycle, transport distance and terminal utilisation. NRG4NOW does not assume that the theoretical -600 to +600 EUR/MWh market range can be captured for the same energy volume; it treats observed spreads as evidence that flexibility has acquired an independent economic value. [35–38]

Design item	Reference requirement
Site	Existing industrial brownfield with verified land, access, safety distances, cooling route and permitting authority
Energy input	Pathway A electricity surplus, Pathway B stranded molecules, or a metered combination with separate ledgers
Storage fleet	Multiple qualified OEM modules; no mandatory TSTM geometry or supplier
Conversion block	100 MW net continuous target plus project-defined redundancy; OEM efficiency and auxiliary demand entered explicitly
Reserve duration	Declared in hours or days from usable inventory, leakage, conversion availability and replenishment time—not labeled seasonal without evidence
Acceptance test	Metered net output, emissions, leakage, water, acoustic, availability, failover and cost per useful MWh

7.3 Capital Exposure

The original model combines three distinct quantities: schedule delay, an assumed daily opportunity cost and an assumed value per consumed megawatt-hour. These are not guaranteed revenues. The investor model therefore treats them as scenario variables and reports the arithmetic separately from the interpretation.

Scenario	Delay	Daily opportunity cost	Energy-value assumption	Indicative exposure [18, 19]
Conservative	36 months	USD 75,000/day	USD 75/MWh	Schedule and energy components calculated separately
Reference	60 months	USD 150,000/day	USD 110/MWh	Central planning case; requires workload evidence
High-impact	76 months	USD 250,000/day	USD 150/MWh	Original upper case; not a forecast or guaranteed token margin

In the original high-impact case, 76 months × 30 days × USD 250,000/day is approximately USD 570 million. Separately, 100 MW × 8,760 hours/year × 6.3 years × USD 150/MWh is approximately USD 828 million; adding a USD 150 million grid contribution produces USD 1.548 billion. The mathematics is transparent, but the commercial meaning depends on utilization, power-usage effectiveness, chip availability, model-training economics, contracted compute revenue and whether the grid contribution is actually incremental.

7.4 Chip Obsolescence

Hyperscaler economics are governed by a 12–18 month semiconductor innovation cycle, not by a 20–40 year utility-asset cycle. A 60–96 month interconnection delay spans three to five GPU generations. Hardware purchased for hundreds of millions of dollars can become technically and economically obsolete before it processes a single production token. [6]

13. Rapid innovation cycle: leading chip suppliers release a new accelerator generation every 12–18 months, improving performance per watt and training throughput. [6]
14. Stranded-silicon collapse: during a six-to-eight-year grid wait, the market advances through several complete generations while installed racks remain unproductive. [6]
15. Standby devaluation: after 24–36 months without power, older architectures can become too inefficient relative to current systems to support the original business case. [6]

Time after purchase	Conventional grid path	NRG4NOW below 20 months [6]
Month 0–18 / Generation 1	Queue; hardware unproductive	Full operation and maximum token margin
Month 19–36 / Generation 2	No connection; chips overtaken	Investment earning; upgrade funded from cash flow
Month 37–76+ / Generation 3+	Obsolete hardware finally connects	Continuous modular operation with current chips

7.5 Frontier Model Training

The primary objective is not passive operation of existing workloads. It is the immediate training of the next frontier and AGI models on the newest available hardware. A six-to-twelve-month training advantage can determine model leadership, developer adoption, enterprise contracts and the ability to set the next platform standard.

16. Time-window monopoly: frontier training requires tens of thousands of current accelerators operating at uninterrupted 24/7 load. Delayed energy destroys the strategic lead before the model is trained.
17. Chip-vendor sales pressure: semiconductor companies pre-finance advanced fabrication capacity and must deliver and monetize each leading generation before the next tape-out supersedes it.
18. Location independence: net-decoupled molecular logistics allows temporary or permanent training clusters to be deployed at liquid energy hubs without waiting for a conventional site and grid corridor.

The result is direct: the newest chip generation reaches full load immediately after delivery, the model is trained before the hardware becomes obsolete, and the manufacturer can place the next generation without being blocked by the previous customer's grid queue.

7.6 Modular Expansion Versus Monolithic Power Projects

A monolithic power project concentrates schedule, fuel, grid, permitting and technology risk in one site. Modular deployment can add qualified storage, conversion and compute blocks incrementally, substitute suppliers and reroute inventory. The benefit is optionality, not automatic cost superiority: series economics must be demonstrated with quotations, utilization, reserve requirements, transport cycles and acceptance-test data.

CHAPTER 8: GOVERNANCE, COALITION AND PHYSICAL ESG

8.1 Coalition of Winners

A new infrastructure architecture fails when it merely attacks the values and business models of established participants. NRG4NOW instead converts the dominant pain points of five sectors into a shared value engine: the Coalition of Winners.

Sector / actor	Current constraint	NRG4NOW winner state [2, 6, 23–28]
Hyperscalers	60–96 month grid queue; 500+ local bans	Model target below 20 months; reduced grid dependence; site-specific water and acoustic compliance
Ports and terminals	Passive infrastructure beside curtailed energy	Autonomous energy hubs and P2P molecular trade
Infrastructure funds	Long-life immobile corridors and regulatory exposure	Modular hub finance; scenario-based return range; reroutable assets
Offshore producers	$400M write-downs; flaring and stranded gas	Day-one monetization without billion-dollar pipelines
OEMs and logistics	Bespoke assets and margin pressure	Series production of standardized modules and permanent rotation demand

8.1.1 The Five-Pillar Alliance Matrix

19. Hyperscalers and AI compute operators can pursue earlier access to firm power, protect GPU investments and address the physical conditions behind community and Texas audit requirements, subject to permits and project execution. [4, 6]
20. Ports and maritime terminals become Giga-Energy-Hubs that convert low-value or curtailed electricity into transportable molecular inventory.
21. Infrastructure funds finance repeatable modules instead of waiting decades for a single immobile corridor; rerouting limits stranded-asset exposure.
22. Offshore producers and national oil and gas companies monetize fields, associated gas and NGL streams that cannot support a conventional pipeline or LNG megaproject. [2, 16]
23. Industrial OEMs and logistics fleets replace bespoke one-off engineering with Ford-style serial production and closed-loop utilization across road, rail and maritime transport. [23–28]

8.2 From Paper Certificates to Hourly Auditable Physics

The legacy system for certifying sustainable energy relies on annualized net-accounting and unbundled paper transactions (e.g., Renewable Energy Certificates / RECs). Under this system, a data center can legally claim to be "100% green" by purchasing certificates from a solar farm operating in July, while the data center is physically powered by 100% coal generation from the local grid during a January night. [17]

This accounting mechanism is facing an existential regulatory and social backlash in 2026, as the public increasingly views it as corporate greenwashing.

Physical ESG replaces annual paper balancing with hourly, physically auditable molecular tracking:

24. Unbroken Provenance: Through the standardized Module Identity dataset (Patent DE 10 2026 003 523.6), every molecular energy container carries an immutable cryptographic history of its origin, including the exact compression location, primary energy source, timestamp, and CO_2 footprint. [1]

24/7 physical matching evidence: The module ledger can record when a qualified molecule is docked, verified and converted while the data center executes its workload. This creates auditable evidence for a future certification or assurance process but does not by itself establish regulatory sustainability status. [1]

Auditable physical evidence: Module identity and handshake logs can document provenance, operating state, transfer and conversion. These records can support independent audits but do not replace applicable energy, emissions, sustainability or accounting rules. [1]

8.3 Maritime Regulatory Framework

For the maritime deployment of the NRG4NOW Framework (such as modular LPG/NGL extraction in the Orange Basin or CNG shuttles in the Niger Delta), the architecture relies on established international classification societies. [11, 12]

The system leverages a critical regulatory and structural distinction:

Maritime classification and onshore codes: Integration into a vessel does not automatically disapply terrestrial pressure-equipment law. The controlling code basis must be agreed with the flag administration, classification society, port and competent national authorities for the specific module, vessel, route and operation. Applicable requirements may include the ABS Guide for Vessels Intended to Carry CNG in Bulk, DNV Rules for Classification of Ships, pressure-equipment rules and dangerous-goods law. [11, 12]

Weight and cost optimization: Maritime rules may permit design methods based on fracture mechanics and fatigue assessment, but any wall-thickness, mass or payload advantage must be demonstrated for the selected material, pressure cycle, damage case and class rule. The previously stated 35–40% reduction remains a design scenario until supported by a comparative certified calculation. [11, 12]

8.4 Regulatory Gap Analysis

ABS, DNV or China Classification Society rules can define vessel and pressure-equipment requirements for maritime deployment, but classification does not replace the complete approval chain. Each project must allocate responsibilities for pressure systems, hazardous-goods transport, port operations, environmental impact, emissions, water intake and discharge, electrical safety, cyber security and data-centre operation. [11, 12]

Approval layer	Responsible domain	Open project evidence [11, 12]
Module and vessel classification	ABS, DNV, CCS or equivalent	Design basis, fatigue, pressure cycle, fire and collision cases
Port and transport operation	Port authority and hazardous-goods regulators	Berthing, transfer zone, emergency response and route approval
Energy conversion	National and local environmental / safety authorities	Fuel, emissions, methane leakage, noise and water balance
Compute and data	Electrical, cyber-security and telecom authorities	Island-mode protection, fibre redundancy and data governance

CHAPTER 9: STRATEGIC CARRIER APPLICATIONS

9.1 Desert Hydrogen Without Water

The Tabuk–NEOM region combines strong solar and wind resources with maritime access. Conventional green hydrogen additionally requires desalination, electrolysis, water treatment and export infrastructure. NRG4NOW separates renewable compression from that water-intensive production path. [29]

Renewable electricity compresses a qualified molecular carrier into certified exchangeable modules for multimodal shipment. The working pressure is selected from the carrier, OEM design and applicable approval rather than fixed by the architecture. Compression avoids electrolysis water, while auxiliary cooling, cleaning and site services remain part of the project water balance. [1, 29]

9.2 Destination-Side Methane Pyrolysis

At the destination, methane pyrolysis can separate hydrogen and solid carbon. The hydrogen is produced beside the load; the carbon requires a verified quality, market, storage or disposal pathway. [29]

Desert Hydrogen Without Water means avoiding electrolysis at the desert source, not zero total water or zero carbon. The balance must include methane origin, leakage, process heat, carbon yield and quality, transport and end-of-life treatment. [29]

Sequence	Location / operation	Physical result [29]
1	Tabuk / NEOM renewable generation	Solar and wind power P2P compression
2	Qualified multimodal module shuttle	Compressed molecules transported by ship
3	Destination energy hub	Modules dock directly at compute or industrial load
4	Methane pyrolysis at destination	Hydrogen plus solid carbon; no electrolysis water at the desert source

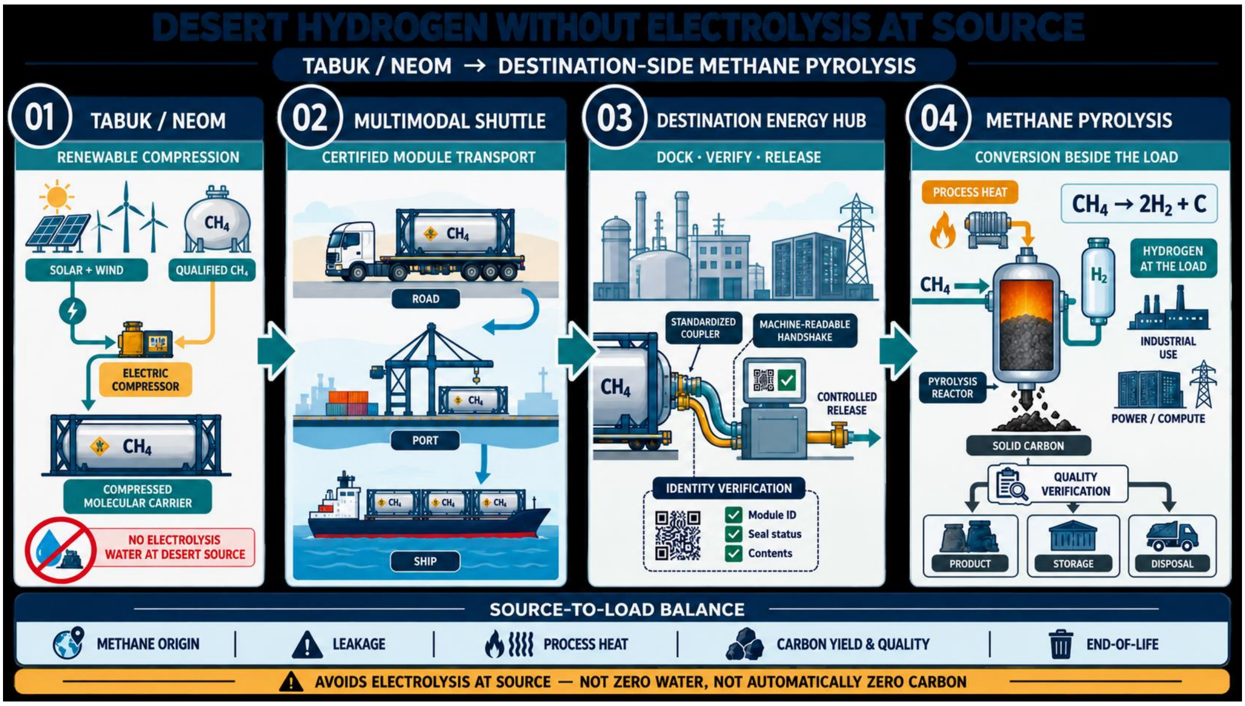

Figure 9.1: *Desert Hydrogen Without Electrolysis at Source (Tabuk/NEOM to Destination Pyrolysis) – Modular logistics sequence showing renewable compression at the desert generation node (01), certified multimodal transport (02), automated docking at the destination hub (03), and decentralized methane pyrolysis beside industrial or compute loads (04), producing hydrogen and sequesterable solid carbon without freshwater consumption for electrolysis at the generation site.*

9.3 Existing Methane Versus Synthetic Methane

Existing or stranded methane and synthetic methane are different energy pathways. Existing methane begins with a physical molecular resource whose origin, leakage and alternative disposition must be recorded. Synthetic methane begins with electricity, a hydrogen source and a certified carbon source and therefore carries additional conversion losses, water and carbon-accounting requirements. Similar downstream modules do not justify one shared efficiency or emissions claim. [29]

9.4 Solid Carbon Accounting and Utilization

Destination-side methane pyrolysis produces hydrogen and solid carbon only when the process and heat source are specified. The carbon is not automatically a permanent sink or saleable product. The project must report carbon yield, purity, energy demand, storage or product route, market capacity and end-of-life treatment before assigning avoided-emissions value. [29]

9.5 China Pilot Corridor

The China application starts with one measured generation cluster, one qualified conversion/compression hub, one transport route and one contracted load centre. Official capacity, Electricity Maps generation and any dispatch record remain separate evidence classes. The corrected Annex China contains the data reconciliation and pilot-gate requirements; this section defines only the architectural link. [20–22]

9.6 Other Carrier-Specific Applications

The same interface logic can be profiled for CNG/RNG, LPG/NGL, hydrogen and other qualified molecular carriers, but each carrier requires its own material, pressure, leakage, transport, conversion and regulatory envelope. Carrier-neutral architecture means common governance and interface logic—not physical interchangeability of gases or vessels. [1, 11, 12]

CHAPTER 10: ROADMAP, CONCLUSION, ACKNOWLEDGEMENTS & REFERENCES

10.1 Conclusion: Availability as a System Architecture

The expansion of renewable generation and AI compute has exposed a mismatch between installed capacity, instantaneous network access and energy availability over time. Transmission, batteries, firm generation and molecular storage solve different parts of this problem. The NRG4NOW proposition is to connect them through a standardized logistics and control architecture rather than to present one technology as a universal substitute.

The strongest immediate use cases are measurable: curtailed or non-connectable generation, associated or stranded molecular streams, brownfield industrial sites, pressure- or capacity-constrained offshore networks and compute projects whose hardware life cycle is shorter than the grid-connection queue. Each case can be tested through an auditable energy balance, a carrier-specific emissions boundary, a permitting map and a time-to-energy sensitivity model. [2, 4, 6, 16, 17, 33]

The White Book therefore distinguishes established facts, observed datasets, scenario quantities and engineering design targets. NRG4NOW becomes investable only when pilot projects demonstrate net deliverable energy, safety, leakage performance, water and acoustic compliance, operating availability and cost per useful megawatt-hour across the complete chain.

10.2 NRG4NOW as the Strategic Umbrella

The NRG4NOW coordination framework is conceived as the roof strategy for NRG4NOW deployment. It is not a generating company and does not replace grid operators, producers, ports, OEMs or hyperscalers. Its purpose is to make energy available by coordinating assets that already exist or can be produced by participating industries. [1]

The alliance is structured around five operating functions:

25. Collect curtailed, pre-emptively throttled or non-connectable renewable electricity at the generation node. [1]
26. Use that electricity to fill carrier-specific molecular storage modules through controlled compression or another qualified conversion path. [1]
27. Move certified modules through road, rail, inland-waterway and maritime rotation chains. [1]
28. Operate energy hubs for inspection, filling, dispatch, conversion, provenance and reserve management. [1]
29. Provide contracted supply security through diversified inventories, routes and conversion modules rather than through a single grid connection. [1]

NRG4NOW supplies the governance layer for the Coalition of Winners: hyperscalers define firm-load requirements; producers supply certified molecules; ports and terminals operate transfer nodes; infrastructure funds finance repeatable assets; and OEMs manufacture interoperable modules. Performance guarantees remain project-specific and depend on inventory, redundancy, permits, carrier balance and contracted operating reserves. [1, 23–28]

Governance domain	Required rule
Ownership and licensing	Patent and specification rights remain explicit; access terms are published and do not imply transfer of ownership
Voting	Balanced representation across demand, supply, logistics, finance and OEM pillars; no single company controls a decision class
Conformance	Public test profiles, accredited laboratories, version control and equal cross-vendor substitution tests
Liability and insurance	Responsibility assigned for module manufacture, filling, transport, docking, conversion and orchestration
Cybersecurity and data	Minimum telemetry, access control, audit retention, incident response and commercially protected data boundaries
Energy and emissions claims	Common source-to-load ledger covering electricity, molecules, leakage, auxiliaries, water and useful output
Appeals and revocation	Independent technical appeal, corrective-action period and transparent suspension or certificate-revocation procedure

10.3 Global Development Roadmap (2026–2028)

30. Q3 2026: Formal publication of the NRG4NOW reference specification on Zenodo and constitution of the international NRG4NOW working groups.
31. Q4 2026: Acquisition of formal Approval in Principle (AiP) for standard Giga-Hubs and TSTM- M containment classes from ABS and DNV. [11, 12]
32. Q2 2027: Commissioning of the Phase 1 reference hubs in the potential configuration of the Hamburg-Moorburg reference scenario (50 modules × 13.28 GWh = 664 GWh gross methane storage capacity at 250 bar; subject to independent validation regarding tanks, gas quality, leaks, service life, and conversion; the net deliverable amount of electricity is lower and depends on the technology), as well as the initial commissioning of the floating pipeline extension in the African offshore area.. [2, 29]

10.3.1 Twin Pilot Programme and Validation Ladder

Revision 16 replaces a direct jump from concept definition to a 500 MW cluster with two linked pilots. The first pilot validates a measured renewable-surplus-to-compression corridor. The second validates a molecular-source-to-brownfield-load corridor, including docking, custody transfer, conversion and continuous operation. Together they test both energy-input pathways defined in Section 3.7 without requiring one technology or supplier to carry the entire programme. [1, 29]

Gate	Required evidence	Scale decision
1. Concept definition	System boundary, carrier, source, load and evidence ledger	Authorize independent review
2. Independent design review	Mass/energy balance, safety case, interfaces and failure modes	Authorize component qualification
3. Component qualification	Pressure, materials, valves, metering, controls and digital identity	Authorize measured corridor pilots
4A. Renewable corridor pilot	Measured electricity input, compression, inventory, transport and delivered output	Validate surplus-energy pathway
4B. Molecular corridor pilot	Measured source gas, loading, transport, docking and brownfield conversion	Validate stranded-resource pathway
5. Integrated hub pilot	Multi-module operation, reserve policy, OEM interoperability and incident response	Authorize replicable deployment
6. Commercial replication	Verified unit economics, permits, contracts and availability	Consider 100–500 MW clusters

Acknowledgements

Special thanks are extended to Dave Stenning for the valuable technical exchange, his reconstruction of the commercial and institutional obstacles that stopped earlier marine-CNG projects, and his decades of pioneering work on maritime-classified pressure-vessel and natural-gas transport systems. [15]

Appendix A — Evidence-Status Register

Statement class	Status in this revision	Publication requirement
Official statistic	China end-2025 PV and wind capacity dated to the NBS release	Archive source and retain reporting boundary
Observed telemetry	August 2026 Electricity Maps/Windy comparison	Preserve timestamp, zone, source field and original screenshot
Journalistic dataset	Data-center rejections and market reports	State method, retrieval date and non-census character
Scenario	Queue energy, 664 GWh Moorburg inventory and opportunity cost	Expose assumptions and do not label as measured result
Design target	20-month deployment, water, acoustic and net-decoupled operation	Convert to verified result only after pilot acceptance testing

Appendix B — 100 MW Energy-Balance Template

Ledger item	Unit	Project entry
Net facility load and IT load	MW	Define boundary and PUE
Daily / annual useful electricity	MWh / GWh	2,400 MWh/day and 876 GWh/year at 100 MW continuous
Gross carrier input	mass and MWh LHV/HHV	From certified composition and metering
Compression / conditioning / transport auxiliaries	MWh	Measured by pathway and route
Conversion efficiency and net output	% and MWh	OEM curve plus site acceptance test
Usable reserve	hours or days	After residual inventory, outages and replenishment time

Appendix C — Vendor-Neutral Module Qualification Basis

Each candidate module enters a common qualification file covering carrier/material compatibility, pressure design, usable inventory, leakage, cycle and fatigue evidence, valves and isolation, fire and impact cases, inspection, transport classification, telemetry, cybersecurity, connector profile and handshake interoperability. Qualification applies to a defined carrier and operating envelope; it is not automatically transferable to another gas, pressure or transport mode. [1, 11, 12]

Appendix D — Stage-Gate and Permitting Map

Gate	Decision evidence	Result
G0 — use case	Site, 100 MW boundary, pathway A/B and commercial owner	Proceed to concept or stop
G1 — resource	Metered energy/molecule availability and ownership	Bankable source envelope
G2 — technology	Qualified OEM modules, compression and conversion package	Frozen reference design
G3 — safety and permits	HAZID/HAZOP, class, pressure, transport, environmental and site approvals	Authority pathway accepted
G4 — finance and contracts	CAPEX/OPEX, insurance, supply, offtake and contingency	Financial close
G5 — build and integration	Factory tests, civil works, interface and cyber tests	Commissioning release
G6 — acceptance	Net energy, availability, leakage, emissions, water, acoustics and failover	Operate, correct or reject

Appendix E — 24/7 Power-Supply Profiles in Major U.S. Data-Centre Regions

E.1 Scope and Time-Zone Method

This appendix preserves the time-stamped Electricity Maps observations used in Sections 2.6.1 and 2.6.2. All interface timestamps are shown in GMT+2 and converted to the applicable local daylight-saving time on 20 August 2026: PDT for CAISO, CDT for ERCOT and EDT for PJM. The figures are observational snapshots of regional generation; they are not plant-level dispatch records and do not identify the electricity supply of any individual data centre. [41]

E.2 CAISO: Four Phases Across the Solar Cycle

The four CAISO observations span the pre-sunrise period, local solar noon, the post-sunset evening ramp and midnight. Together they show that the gas fleet remains active during every observed phase, with the lowest absolute output at noon and the highest observed output after sunset. [41]

Interface timestamp	Converted local time	Observed phase	Gas output	Gas share
20 Aug, 14:00 GMT+2	20 Aug, 05:00 PDT	Before sunrise	13.4 GW	49.47%
20 Aug, 21:00 GMT+2	20 Aug, 12:00 PDT	Solar noon	11.5 GW	27.55%
21 Aug, 04:15 GMT+2	20 Aug, 19:15 PDT	Post-sunset evening	17.8 GW	40.21%
21 Aug, 09:00 GMT+2	21 Aug, 00:00 PDT	Midnight	15.4 GW	49.59%

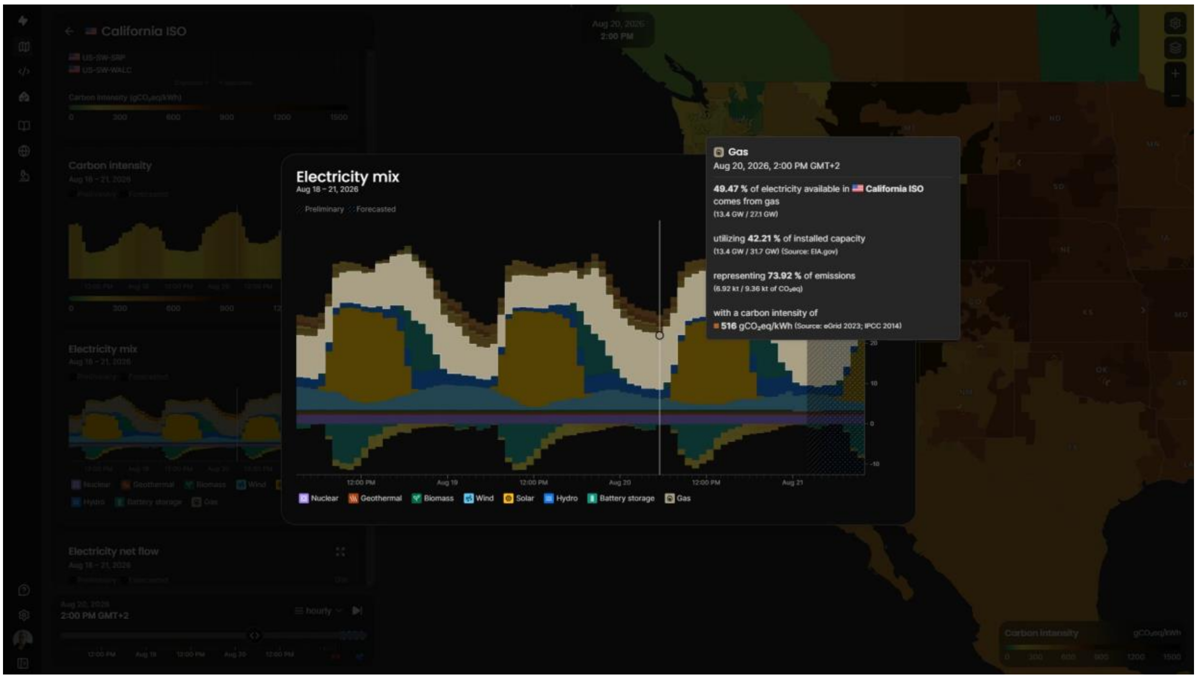

Figure E.1: CAISO before sunrise, 20 August 2026, 05:00 PDT — 13.4 GW gas generation and a 49.47% gas share. [41]

Figure E.2: CAISO at local noon, 20 August 2026, 12:00 PDT — 11.5 GW gas generation and a 27.55% gas share. [41]

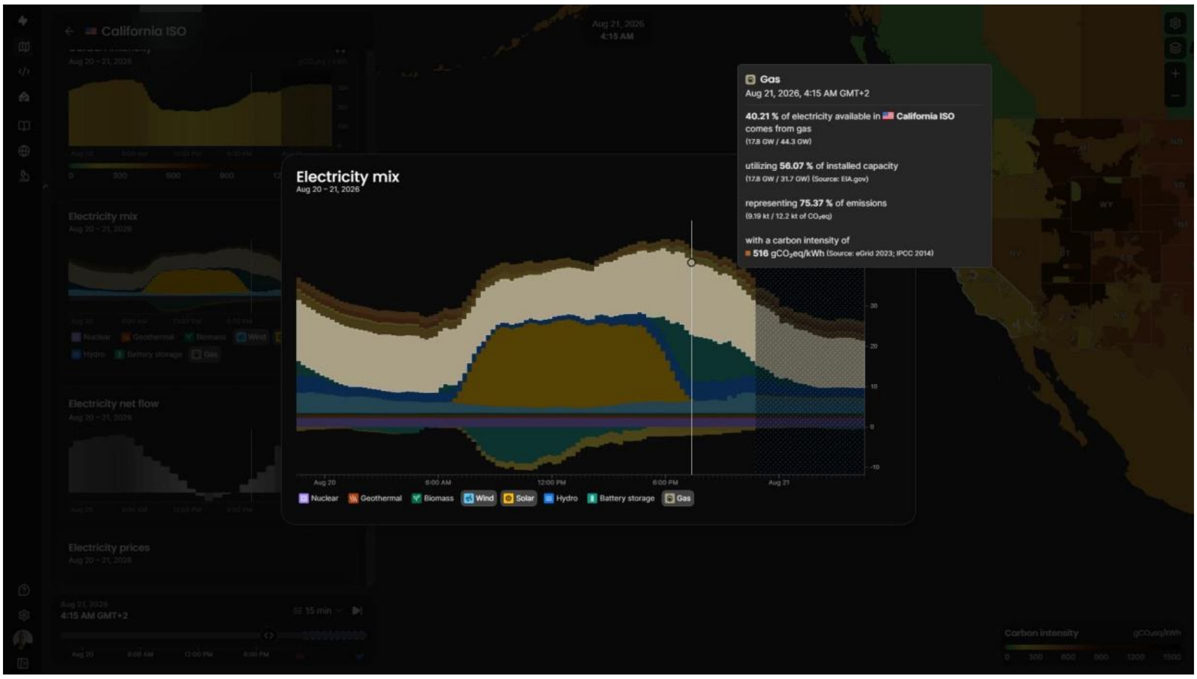

Figure E.3: CAISO after sunset, 20 August 2026, 19:15 PDT — 17.8 GW gas generation and a 40.21% gas share. [41]

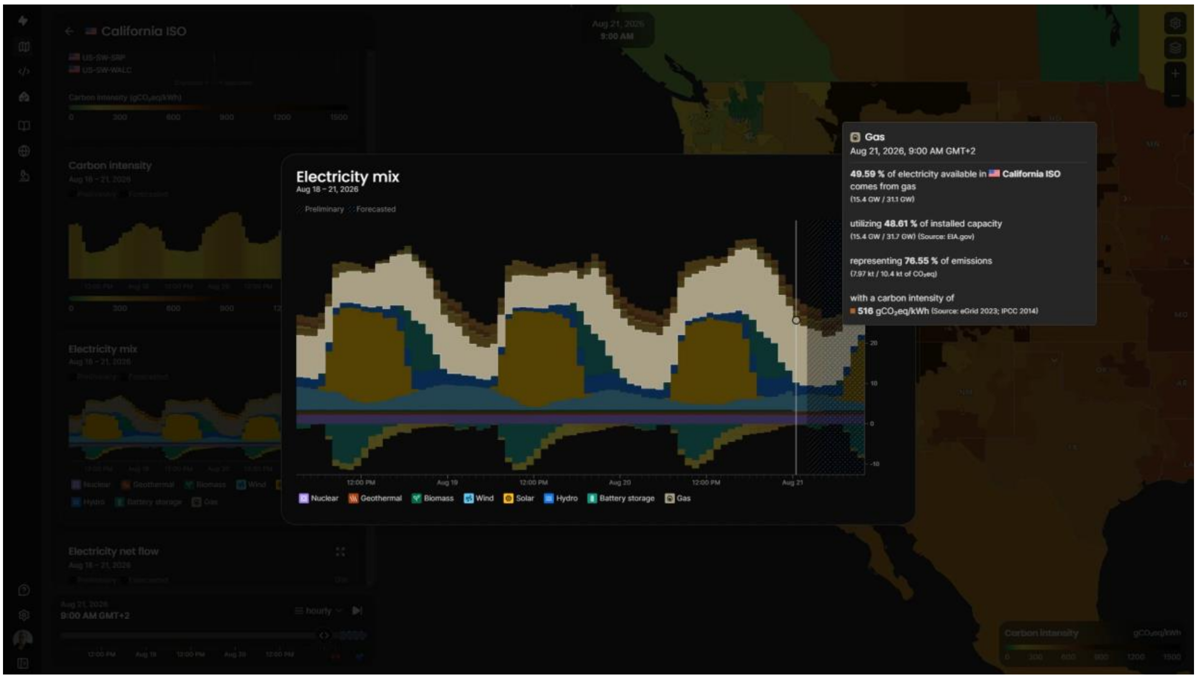

Figure E.4: CAISO at midnight, 21 August 2026, 00:00 PDT — 15.4 GW gas generation and a 49.59% gas share. [41]

E.3 Matched Noon-to-Evening Comparison

The following pairs compare each regional system at local noon and after sunset on the same local date. CAISO and ERCOT show large evening increases in gas-fired output. PJM remains strongly gas-dependent at both observation times. [41]

Figure E.5: ERCOT at local noon, 20 August 2026, 12:00 CDT — 30.3 GW gas generation and a 35.58% gas share. [41]

Figure E.6: ERCOT after sunset, 20 August 2026, 19:00 CDT — 44.5 GW gas generation and a 50.74% gas share. [41]

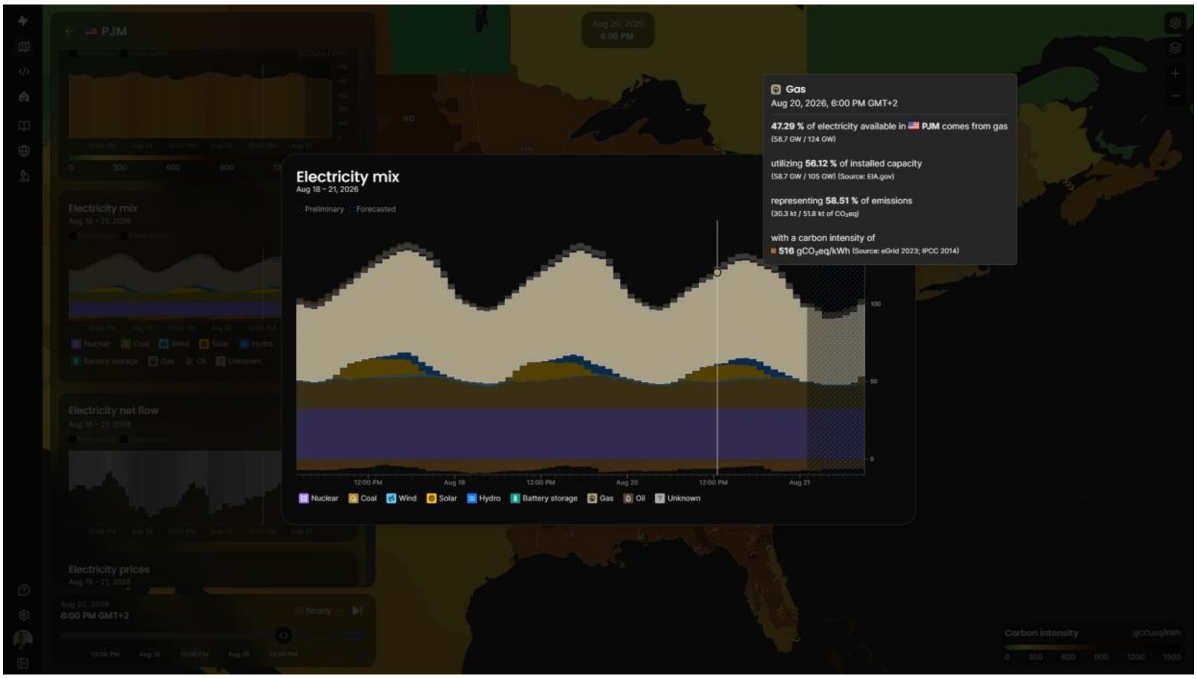

Figure E.7: PJM at local noon, 20 August 2026, 12:00 EDT — 58.7 GW gas generation and a 47.29% gas share. [41]

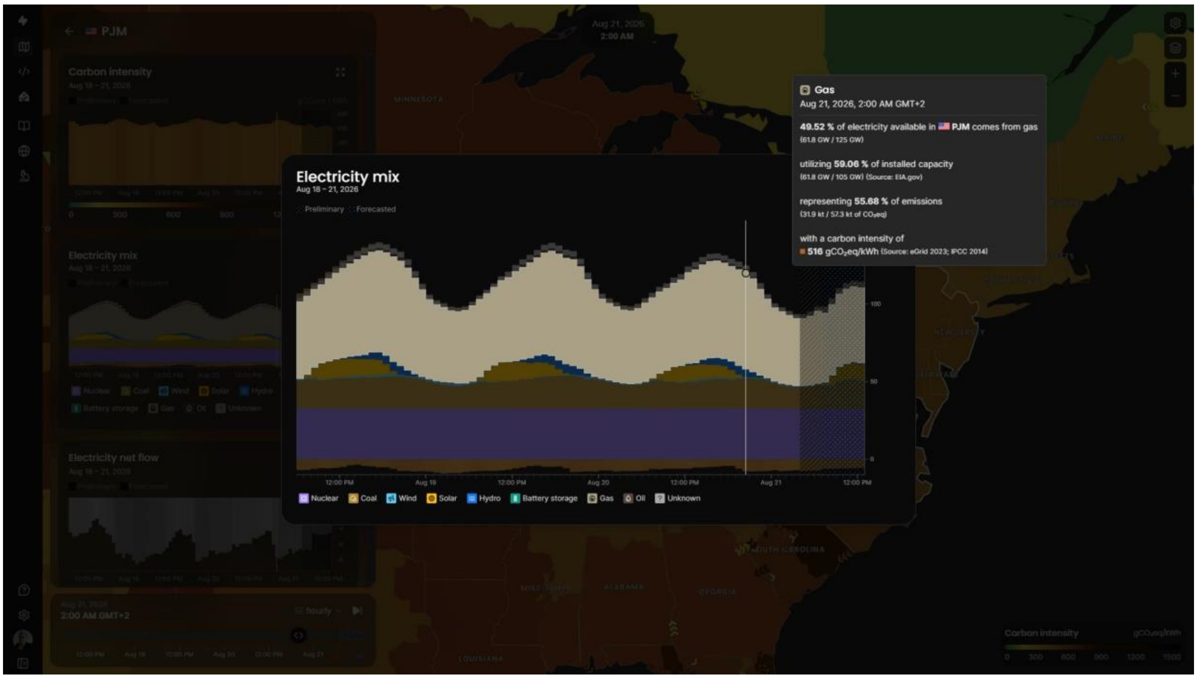

Figure E.8: PJM after sunset, 20 August 2026, 20:00 EDT — 61.8 GW gas generation and a 49.52% gas share. [41]

E.4 Interpretation Boundary

The screenshots directly establish the reported regional generation values, technology shares and timestamps displayed by Electricity Maps. They support comparison of the daytime and evening supply structures. They do not establish fuel consumption in volumetric units, assign power flows to individual data centres, identify contractual supply arrangements or prove a single causal explanation for generator dispatch. The term gas generation therefore refers to electrical output attributed to gas-fired generation, not to measured physical gas consumption. [41]

Reference Register

1. Dzikowski, R. (22 June 2026). Standardisierte Energie-Logistikarchitektur für molekulare Energieträger mit austauschbaren Speicher-, Transport-, Umwandlungs-, Netz- und Hub-Modulen. Patent Application DE 10 2026 003 523.6, DPMA.
2. Dzikowski, R. (May 2026). The Architecture Gap: Structural Offshore Gas Monetisation Constraints Across Five Producer Regions. Zenodo Repository.
3. Taft, M. (4 August 2026). How Data Centers Broke American Politics. WIRED Magazine.
4. Walton, R. (5 August 2026). Facing an estimated 474 GW of interconnection requests, Texas hits pause on data centers. Utility Dive.
5. Bryce, R. (August 2026). Data Center Rejection Database and 2026 rejection reports.
6. The Information / New York Post Business Report (10 August 2026). More than 500 US jurisdictions blocking data-center construction.
7. Reuters Events / Grid Flexibility (4–5 August 2026). California curtailment reached 4.5 TWh in the first half of 2026.
8. Electricity Maps, Fraunhofer ISE, SMARD, REN, IPTO/ADMIE, ERCOT and Windy. Operational telemetry, 3–13 August 2026 and January 2026.
9. Financial Times, Adam Tooze (6 June 2026). Wasting China's solar panel surplus is madness.
10. Global Wind Energy Council / BCG. Wind supply-chain and industrial capacity reports.
11. American Bureau of Shipping. Guide for Vessels Intended to Carry CNG in Bulk.
12. Det Norske Veritas. Rules for Classification of Ships, Part 5 – CNG Carriers.
13. World Bank Group / Tractebel Engineering S.A. (2015). CNG for Commercialization of Small Volumes of Associated Gas. Final public main report, Document 104200-V2-WP, Version 2. Global Gas Flaring Reduction programme. The report's evidence base and project-status statements reflect information available in 2014–2015.
14. Loebs, B. (15 October 2025). World's first CNG vessel delivers major cost savings for PLN. Petromindo. Reports JAYANTI BARUNA operations between Gresik and Lombok, first gas-in in September 2024, 25 voyages, delivered gas, vessel capacity and PLN cost comparisons.
15. Private technical correspondence with a veteran marine-CNG developer (24–26 March 2026). Historical development of bottle ships, Coselle and Optimum; ABS classification; manufacturing, financing and project-interface barriers. Author's archive.
16. Nigerian Upstream Petroleum Regulatory Commission. Annual Statistical Bulletin 2024. Abuja, 2025.
17. BloombergNEF / Bloomberg (May 2026). Europe solar curtailment, negative prices and grid investment requirements.
18. Dzikowski, R. (June 2026). Copilot Antworten 1–4: AI-generated strategic, investor, engineering and standardization stress tests for NRG4NOW. Development documentation; simulated scenarios, not Microsoft corporate memoranda.
19. Dzikowski, R. (June 2026). Meta Antworten 1–4: AI-generated hyperscaler, investor, engineering and institutional stress tests for NRG4NOW. Development documentation; simulated scenarios, not Meta corporate memoranda.
20. National Bureau of Statistics of China (28 February 2026). Statistical Communiqué of the People's Republic of China on the 2025 National Economic and Social Development; end-2025 grid-connected wind and solar capacity.
21. National Energy Administration of China (22 July 2026). National electricity statistics for January–June 2026; installed solar capacity approximately 1.27 TW and wind capacity approximately 0.68 TW at 30 June 2026.
22. Electricity Maps (2025 annual view), based on Ember yearly electricity data. China synthetic generation dataset: approximately 1,180 TWh solar, 1,130 TWh wind and 10,400 TWh total electricity; screenshots retained in the empirical evidence file.
23. Hexagon Agility. Mobile Pipeline TITAN and X-STORE CNG/RNG transport modules; official product and 2025 order information.
24. NPROXX. Type 4 hydrogen pressure vessels and certified 500-bar MEGC transport system; official product and certification information.
25. Luxfer Gas Cylinders. G-Stor Hydrosphere MEGC and Type 3/4 hydrogen containment systems; official product information.

26. CIMC Enric. CNG tube skids and MEGC products; official product information.

27. Ariel Corporation; Burckhardt Compression; HOERBIGER; Everllence. Official compressor product and partnership information for natural gas, RNG, hydrogen and maritime applications.

28. INNIO Jenbacher; Siemens Energy; Bloom Energy. Official information on gas-engine, gas-turbine and solid-oxide-fuel-cell conversion systems.

29. Dzikowski, R. (2025). Addressing Global Energy Transition Challenges with Modular TSTM Solutions 2.1. Zenodo. DOI: 10.5281/zenodo.15586051. Concept-development reference for TSTM-M and Subsea Access Nodes.

30. The Telegraph. UK wind-curtailment reporting cited in Section 2.7: £960 million, described as almost £1 billion. The publication date and article metadata should be retained with the archived source copy.

31. Dzikowski, R. (2026). Time-stamped Electricity Maps observation log, Germany, West Denmark and ERCOT, 14–17 August 2026. Author screenshots; observational evidence, not a direct curtailment dataset.

32. ntv Nachrichten (10 August 2026). Report on Amazon financing the 7.65 GW 'GW Ranch' off-grid gas-power project in Texas, comprising 35 turbines for an AI data-centre cluster. Archived source cited in the development dossier.

33. Lawrence Berkeley National Laboratory (2025). Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection. US interconnection-queue dataset and technology composition used in Sections 1.2, 2.3 and 7.1.

34. Deutscher Bundestag (23 October 2023). Antwort der Bundesregierung auf die Kleine Anfrage: Abregelung erneuerbarer Stromerzeugung und Anpassung des Strommarktdesigns. Bundestagsdrucksache 20/9016, especially pp. 2-3.

35. Green Energy Tools, Wittstock, L. (updated 26 June 2026). Negative Strompreise Deutschland 2026: Aktuelle Statistik & Live-Tracker. ENTSO-E Transparency Platform, DE-LU Day-Ahead Sequence 1 data in 15-minute resolution; methodology and forecast status retained.

36. EPEX SPOT (April 2026). SDAC communication note: Harmonised Minimum Clearing Price set to -600 EUR/MWh from 28 April 2026 after the previous -500 EUR/MWh floor was reached on 25 April 2026.

37. LevelTen Energy (23 July 2026). European Corporates Choose Hybrid PPAs Over Solar, As Co-location Becomes Essential. Q2 2026 European PPA Price Index commentary; 29% figure limited to the offers represented in the LevelTen dataset.

38. Brown, D. (17 August 2026). Balance of Power commentary on European negative prices, price spreads, hybrid PPAs and the migration of value from generation-only assets to flexibility. LinkedIn post; interpretive market commentary.

39. Sun Tzu. The Art of War, Chapter VII, 'Maneuvering'. Translated by Lionel Giles, 1910. MIT Classics Archive, https://classics.mit.edu/Tzu/artwar.html.

40. US Energy Information Administration (16 June 2026). Solar and natural gas generation in CAISO; utility-scale solar capacity reached 25 GW in April 2026. EIA, Today in Energy. EIA also reports that California held approximately 34% of total US small-scale solar PV capacity at the end of 2025 and remained the country's largest distributed-PV state.

41. Dzikowski, R. (20–21 August 2026). Time-stamped Electricity Maps observation log for CAISO, ERCOT and PJM: matched local-noon and post-sunset generation profiles, including four CAISO solar-cycle phases. Author screenshots; regional observational evidence, not plant-level dispatch or data-centre allocation records.

ANNEX CHINA

SEASONAL MOLECULAR BUFFERING AS A COMPLEMENT TO THE NATIONAL GRID AND INDUSTRIAL POLICY

Corrected evidence and system-analysis annex to the NRG4NOW Global Energy Logistics White Book

Author: Ryszard Dzikowski
DESERTSOLAR UG · Berlin, Germany

Patent Application: DE 10 2026 003 523.6
German Patent and Trade Mark Office (DPMA), 22 June 2026

1 EXECUTIVE FINDING AND EVIDENCE BOUNDARY

CORE FINDING China's renewable build-out is physically real and historically large. It does not follow, however, that every installed gigawatt is available in every hour or that annual generation data can be divided by a later capacity snapshot without qualification. The central analytical task is to keep capacity, generation, production-potential denominators and curtailment separate.

This annex corrects the earlier China draft and narrows it to claims that can be traced to an identified dataset or expressed transparently as a project hypothesis. It does not use the US–EU interconnection-queue scenario of 11,000 TWh, and it does not convert an unverified national curtailment percentage into a measured TWh total.

The resulting diagnosis is sharper, not weaker: China's capacity scale has advanced faster than the temporal, spatial and commercial mechanisms required to turn every possible renewable kilowatt-hour into useful energy at the required place and time. NRG4NOW is presented as a complementary pilot architecture for qualified surplus or otherwise unusable energy—not as a substitute for UHVDC, batteries, pumped hydro, coal, hydrogen policy or the national grid.

Evidence classes used

Evidence class	Meaning in this annex
Official statistic	Dated national source with reporting date and boundary.
Dataset observation	A value displayed by Electricity Maps and retained in the supplied screenshot.
Derived value	Transparent arithmetic from stated inputs; not a new measurement.
Scenario / hypothesis	A planning input requiring site and dispatch validation before investment use.
Design target	A result to be demonstrated in a pilot acceptance test.

2 CHINA'S CAPACITY SEQUENCE: THREE DIFFERENT DATES

A capacity figure is meaningful only with its reporting date. The values below are not competing descriptions of the same instant.

Reporting boundary	Solar	Wind	Status
Electricity Maps 2025 production-potential denominator	≈888 GW implied	≈522 GW implied	Derived from 7,780 / 4,570 TWh potential divided by 8,760 h; broadly an earlier fleet basis
31 Dec 2025	1,201.73 GW	640.01 GW	Official installed capacity, National Bureau of Statistics
30 Jun 2026	1,274 GW	679 GW	Official rounded/precise values, National Energy Administration

Sources: National Bureau of Statistics of China, Statistical Communiqué on 2025 (28 February 2026); National Energy Administration, first-half 2026 renewable-energy and power statistics (22 and 30 July 2026).

The often cited 1,270 GW solar and 680 GW wind figures are therefore correct rounded June-2026 capacity values. Electricity Maps' lower implied 888/522 GW values do not invalidate the 2025 generation data; they show that the application's production-potential denominator uses an earlier capacity basis. This is a time-base mismatch, not evidence that 382 GW of solar equipment disappeared.

3 2025 GENERATION: WHAT ELECTRICITY MAPS ACTUALLY REPORTS

The supplied Electricity Maps annual views identify the following 2025 quantities as synthetic data based on Ember yearly electricity data:

Technology	Generation	Total electricity available	Share	Displayed potential	Potential utilization
Solar	1,180 TWh	10,400 TWh	11.36%	7,780 TWh	15.23%
Wind	1,130 TWh	10,400 TWh	10.84%	4,570 TWh	24.71%
Combined	2,310 TWh	10,400 TWh	22.20%	12,350 TWh	18.70%*

**The combined 18.70% is a derived ratio, 2,310 / 12,350. It is not an Electricity Maps headline value and is not the share of China's total electricity.*

DO NOT MIX THE DENOMINATORS 22.20% is the combined share of reported 2025 electricity available. 18.70% is the combined ratio of reported generation to the application's displayed theoretical production potential. Neither number is a curtailment rate.

Dividing 2025 annual generation by the June-2026 fleet would combine a full-year numerator with a later, larger denominator. That arithmetic may be shown only as a sensitivity, not as the 2025 capacity factor. The fleet expanded substantially throughout 2025 and again during the first half of 2026.

4 ELECTRICITY MAPS EVIDENCE RETAINED

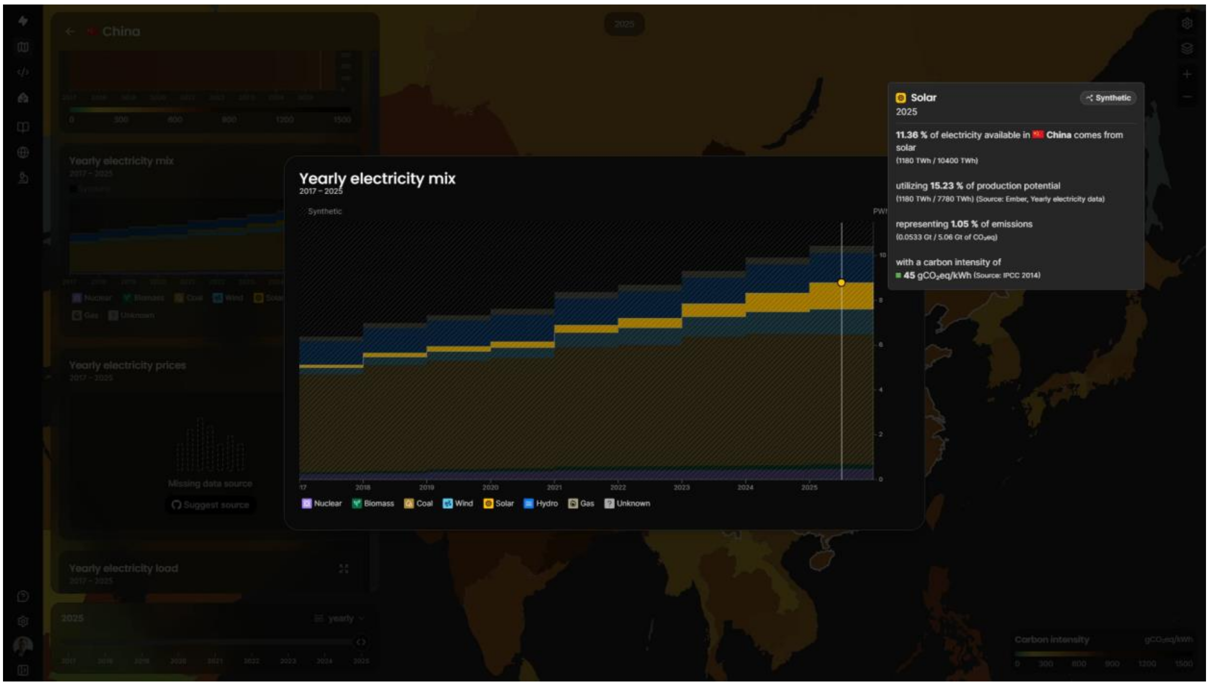

Figure C.1 — China solar generation view for 2025: 1,180 TWh of 10,400 TWh; 11.36%; 7,780 TWh displayed production potential

Evidence status: supplied screenshot; Electricity Maps labels the series 'Synthetic' and cites Ember yearly electricity data. Screenshot retained with the project files.

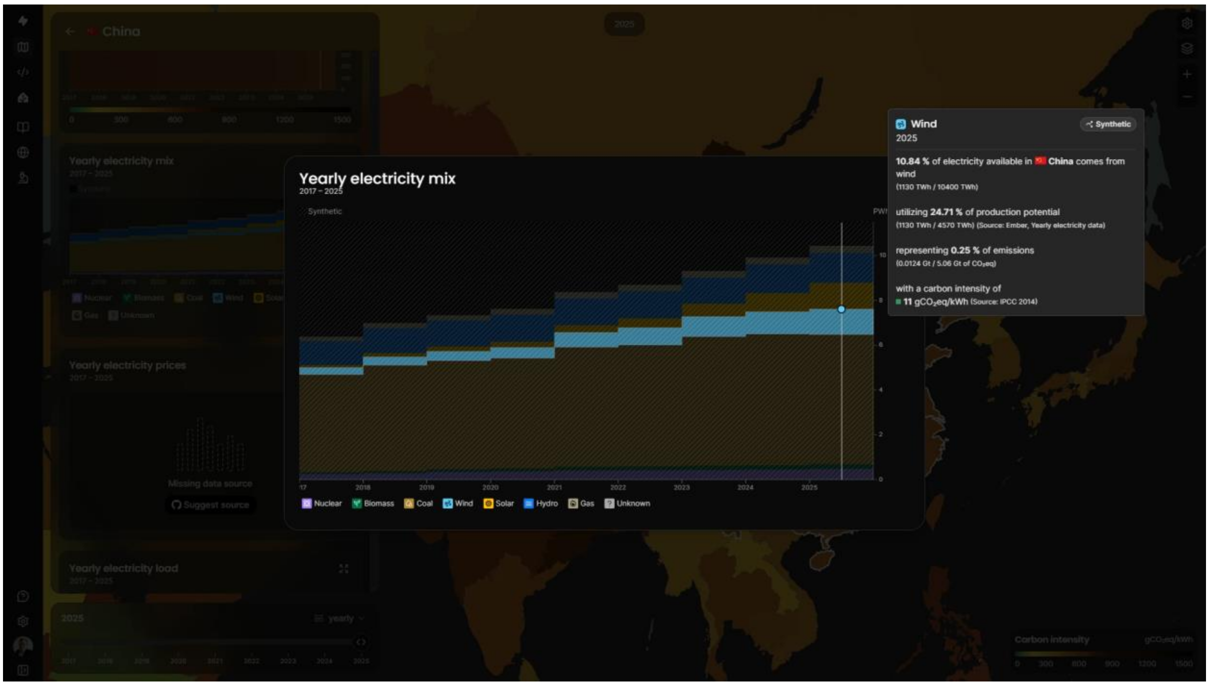

Figure C.2 — China wind generation view for 2025: 1,130 TWh of 10,400 TWh; 10.84%; 4,570 TWh displayed production potential

Evidence status: supplied screenshot; Electricity Maps labels the series 'Synthetic' and cites Ember yearly electricity data. Screenshot retained with the project files.

5 CURTAILMENT: THE BOUNDARY OF WHAT CAN BE CLAIMED

The earlier annex stated 360 TWh curtailed in the first half of 2026 and extrapolated 720 TWh for the full year. Those quantities are removed from the corrected annex because the supplied material does not provide a traceable national dispatch dataset that reproduces them.

A valid national curtailment quantity requires at least four defined terms:

- the reporting boundary: utility-scale only or also distributed generation;
- the counterfactual: available energy, forecast energy, dispatch instruction or technical potential;
- the aggregation method across provinces and market zones;
- the distinction between physical curtailment, delayed grid connection, congestion, commercial non-dispatch and ordinary resource variability.

PUBLICATION RULE No single China-wide TWh curtailment number is used as a measured fact in this annex. A pilot corridor must start from node-level metering and dispatch records. The annex therefore establishes the system question without manufacturing a national loss total.

Observed night-time PV output of zero and low-wind intervals demonstrate intermittency, but they do not demonstrate curtailment. Curtailment exists only where available production is deliberately reduced or cannot be accepted under a defined reporting rule.

6 THE STRUCTURAL DIAGNOSIS: SPATIAL TRANSFER IS NOT SEASONAL STORAGE

China's UHVDC system addresses distance by moving electricity from inland generation regions toward load centres. It is indispensable. A transmission corridor nevertheless transfers power only when generation, network capacity and demand coincide. It cannot carry an August kilowatt-hour into January.

System layer	Primary function	Boundary relevant to this annex
UHVDC and AC grids	Spatial transfer and system interconnection	Do not themselves store energy across seasons
Batteries	Fast response and short-/medium-duration balancing	Economic duration and use case are project-specific; no absolute exclusion is made
Pumped hydro	Bulk electrical storage where geography and permits permit	Site-constrained and not available at every surplus node
Coal and other firm generation	Dispatchable adequacy and stability	Provide energy rather than store a renewable summer surplus
Hydrogen pathways	Carrier conversion, industry and long-duration options	Efficiency, water, infrastructure and end-use boundaries depend on the selected pathway
NRG4NOW pilot layer	Modular molecular capture, storage and logistics	Must prove net energy, leakage, cost, safety and useful destination output

The dual-track strategy—rapid renewable expansion alongside firm generation and major grid construction—is therefore not internally contradictory. It reflects different system functions. The NRG4NOW proposition is to test whether an additional logistics layer can economically capture a defined subset of energy that would otherwise remain unusable at a particular node and time.

7 NRG4NOW PILOT CORRIDOR FOR CHINA

PILOT, NOT NATIONAL EXTRAPOLATION The first decision unit is one measured generation cluster, one qualified conversion/compression hub, one transport route and one contracted demand centre. National-scale claims follow only after replicated acceptance tests.

Operational sequence

Step	Operation	Required evidence
1	Identify a constrained or otherwise unusable energy stream	Metered energy, ownership and dispatch state
2	Select an existing or synthetic molecular carrier pathway	Separate mass, LHV/HHV, carbon origin and conversion ledger
3	Compress and load qualified modules	OEM curves, auxiliaries, leakage and cycle limits
4	Store and rotate above-ground inventory	Usable inventory, residual pressure, transport and safety approvals
5	Deliver to an industrial or compute demand hub	Contracted demand and route availability
6	Convert or consume at destination	Net useful output, emissions, carbon disposition and availability
7	Audit the complete source-to-load chain	Cost, water, energy, leakage and carbon balance

The architecture remains vendor-neutral. Chinese and international manufacturers of pressure modules, compressors, valves, telemetry, engines, turbines, fuel cells and transport equipment may qualify against the same published interface and conformance rules. No specific TSTM vessel geometry is required.

The methane pathway must distinguish two cases. Existing or stranded methane is an energy-resource logistics pathway. Synthetic methane made from renewable electricity is a Power-to-Gas pathway with additional conversion inputs and losses. They may use similar downstream equipment, but they must not share one undifferentiated efficiency or carbon claim.

8 GOVERNANCE, STANDARDS AND INVESTMENT GATES

Domain	China-compatible implementation
Standards	Interface profiles can be mapped to GB/T and relevant SAC processes; maritime and pressure systems require the applicable Chinese and international rules.
Classification	China Classification Society or an accepted equivalent may classify the vessel/module application; class does not replace national permits.
Conformance	Accredited cross-vendor tests for connectors, valves, digital identity, telemetry, isolation and controlled transfer.
Governance	Balanced participation by demand, supply, logistics, finance and OEM actors; patent and licensing rights remain explicit.
Accounting	A common ledger for electricity, molecular inventory, auxiliaries, leakage, water, carbon origin and useful output.

Stage gates

Gate	Minimum decision evidence
G0 — Use case	Named site, commercial owner, energy pathway and destination load
G1 — Resource	Metered source and dispatch/ownership status
G2 — Technology	Qualified OEM modules, compression and conversion package
G3 — Safety and permits	HAZID/HAZOP, class, pressure, transport, fire and environmental pathway
G4 — Finance and contracts	CAPEX/OPEX, insurance, supply, transport, offtake and contingency
G5 — Build and integration	Factory tests, civil works, interfaces, telemetry and cyber tests
G6 — Acceptance	Net energy, availability, leakage, emissions, water, acoustics, failover and cost

9 CONCLUSION

China's official statistics establish the scale: 1,274 GW solar and 679 GW wind were installed by 30 June 2026. The supplied Electricity Maps views establish a different fact for 2025: 1,180 TWh solar and 1,130 TWh wind, together 22.20% of 10,400 TWh electricity available. These statements are compatible once their dates and denominators are kept separate.

The corrected annex does not claim that the difference between theoretical potential and generation is curtailed energy. It does not treat night-time zero output as curtailment. It does not transpose a US–EU queue-energy scenario to China. It does not use an unverified national 360/720-TWh curtailment estimate.

The investment proposition is therefore precise: identify a measurable node where useful energy is being rejected, constrained or stranded; establish ownership; apply a carrier-specific source-to-load balance; qualify vendor-neutral equipment; and demonstrate whether modular above-ground molecular logistics can deliver useful seasonal or long-duration value at lower total system cost and risk than the available alternatives.

DECISION Proceed only to a metered pilot-corridor feasibility study. Do not proceed to a national TWh claim, a fleet order or a bankability statement until G0–G3 evidence has been completed.

10 REFERENCES AND EVIDENCE REGISTER

1. National Bureau of Statistics of China (2026). Statistical Communiqué of the People's Republic of China on the 2025 National Economic and Social Development, 28 February 2026. https://www.stats.gov.cn/english/PressRelease/202602/t20260228_1962661.html
2. National Energy Administration (2026). National power statistics for January–June 2026, 22 July 2026. https://www.nea.gov.cn/20260722/3b678308556b4df0a574537b59a28731/c.html
3. National Energy Administration (2026). Press conference on renewable-energy grid-connected operation in the first half of 2026, 30 July 2026. https://www.nea.gov.cn/20260730/3ce671c387574eeeb120fc3825be0399/c.html
4. Electricity Maps (2026). China yearly electricity-mix views for 2025; synthetic series citing Ember yearly electricity data. Project screenshots Solar_2025.png and Wind_2025.png, supplied August 2026.
5. Dzikowski, R. (2026). NRG4NOW Global Energy Logistics White Book, Revision 12; Sections 2.2.1 and 2.10 and Appendices A–D.
6. Dzikowski, R. Patent Application DE 10 2026 003 523.6, German Patent and Trade Mark Office, 22 June 2026.

Acknowledgement: Dave Stenning is acknowledged for historical technical correspondence concerning compressed-natural-gas transport concepts.

AUGUST 2026 CROSS-SYSTEM OBSERVATION ANNEX

Time-Stamped Electricity Maps and Windy Evidence
3–13 August 2026

Author: Ryszard Dzikowski
DESERTSOLAR UG, Berlin, Germany

Supplementary evidence annex to the NRG4NOW White Book
Revision 1 · August 2026

S.1 Purpose and evidentiary boundary

This annex preserves and compares time-stamped public-interface observations collected between 3 and 13 August 2026. Its purpose is narrow: to document material differences in reported renewable-energy utilization between neighbouring or closely connected electricity systems under broadly comparable temporal and meteorological conditions.

The annex does not derive national curtailment volumes from Electricity Maps or Windy. Curtailment, redispatch and congestion management require separate regulatory, plant-level or system-operator evidence. The screenshots are used as observations that establish a reproducible question: why can neighbouring systems display substantially different ratios between reported output and listed installed capacity at the same or nearly the same time?

S.2 Metric discipline

Displayed metric	What it represents	What it does not prove
Utilizing x% of installed capacity	Instantaneous reported output divided by the platform's listed installed capacity.	It is not the percentage of physical installations switched on.
x% of electricity available	Share of the displayed electricity mix attributed to the selected source.	It is not a capacity factor or a curtailment rate.
Daily production potential	A platform-derived daily comparator used in the 8 August aggregate view.	It is not comparable with a midday instantaneous utilization percentage.
Preliminary / synthetic	A data-status label attached by the platform.	It is not independently verified plant-level telemetry.

S.3 Dataset and selection logic

The supplied archive contained 65 files; the separately recovered 8 August daily screenshot raised the reviewed set to 66. The archive includes country and control-area screenshots, Windy weather maps, multi-day electricity-mix views and several duplicate or contextual captures. The principal evidence set below retains synchronized or near-synchronized country pairs, while controls, duplicate views and screenshots with ambiguous internal timestamps remain documented but are not used as unqualified pairwise proof.

1. Primary comparison group: Germany, East Denmark and Poland.
2. Primary Iberian group: Spain and Portugal.
3. Context controls: Greece, Belgium, the Netherlands, Romania, ERCOT and California ISO.
4. Meteorological context: Windy captures for selected dates; these indicate broad cloud fields but are not site-level irradiance measurements.

S.4 Quantitative overview

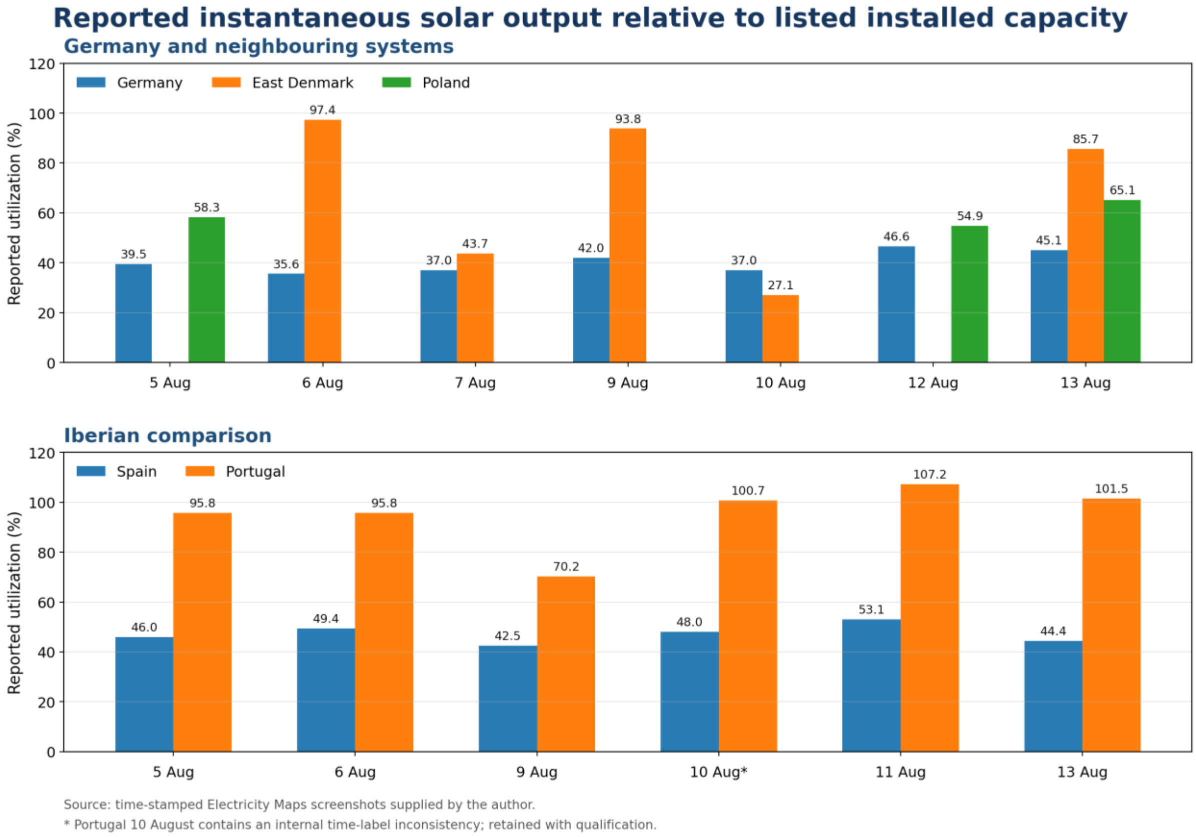

Figure S.1: *Reported instantaneous solar output relative to the installed-capacity value listed by Electricity Maps for the selected synchronized observations.*

Source: author-supplied Electricity Maps screenshots. Missing observations are left blank; no interpolation is used.

Date	Time GMT+2	System A	A (%)	System B	B (%)
5 Aug	12:15–12:30	Germany	39.50	Poland	58.34
6 Aug	12:00	Germany	35.57	East Denmark	97.39
7 Aug	12:00	Germany	37.03	East Denmark	43.72
9 Aug	12:00	Germany	41.96	East Denmark	93.84
10 Aug	12:00	Germany	36.96	East Denmark	27.11
12 Aug	13:30	Germany	46.60	Poland	54.92
13 Aug	12:00	Germany	45.10	East Denmark	85.70
13 Aug	12:00	Germany	45.10	Poland	65.10
5 Aug	12:15	Spain	46.04	Portugal	95.83
6 Aug	12:45	Spain	49.43	Portugal	95.76

Date	Time GMT+2	System A	A (%)	System B	B (%)
9 Aug	12:00	Spain	42.45	Portugal	70.25
10 Aug*	12:00 display	Spain	48.00	Portugal	100.74
11 Aug	13:30	Spain	53.06	Portugal	107.23
13 Aug	12:00	Spain	44.37	Portugal	101.48

* The Portugal screenshot dated 10 August contains a visible inconsistency between the interface display time and the popup's internal local-time line. It is retained as qualified evidence and is not used as a precision timing benchmark.

S.5 Germany, East Denmark and Poland

The northern comparison does not produce one fixed hierarchy. East Denmark substantially exceeds Germany on 6, 9 and 13 August, while the two values are much closer on 7 August and East Denmark falls below Germany on 10 August. Poland exceeds Germany in each retained same-day comparison. This variability is the finding: geographic proximity and a common clock do not force equal reported utilization.

GERMANY · 5 AUG · 12:30 GMT+2 · 39.50% POLAND · 5 AUG · 12:15 GMT+2 · 58.34%

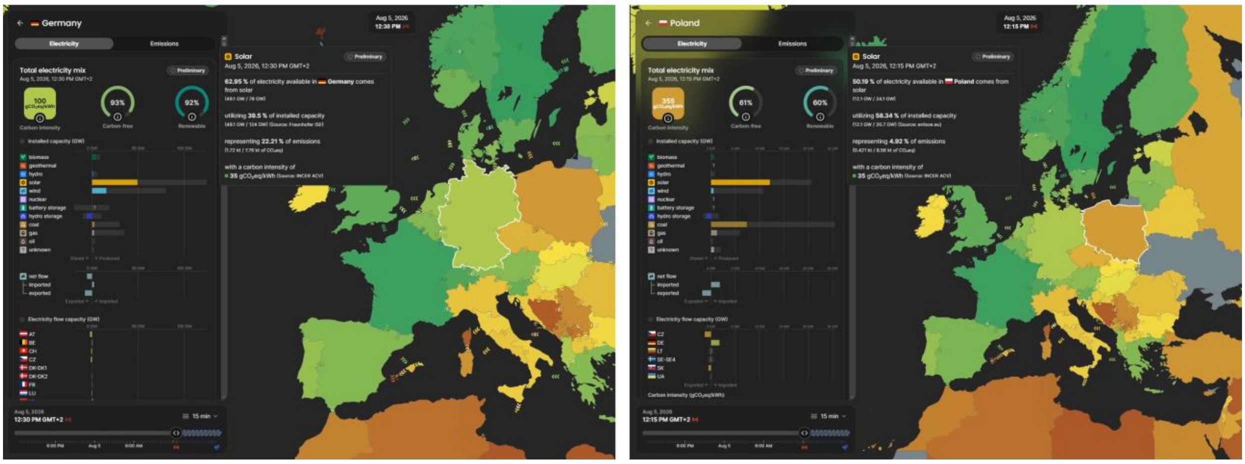

Figure S.2: *Germany and Poland on 5 August. Poland's reported utilization is 18.84 percentage points higher; the captures differ by 15 minutes.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

GERMANY · 6 AUG · 12:00 GMT+2 · 35.57% **EAST DENMARK · 6 AUG · 12:00 GMT+2 · 97.39%**

Figure S.3: *Germany and East Denmark at 12:00 GMT+2 on 6 August. The reported utilization gap is 61.82 percentage points.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

GERMANY · 7 AUG · 12:00 GMT+2 · 37.03%

EAST DENMARK · 7 AUG · 12:00 GMT+2 · 43.72%

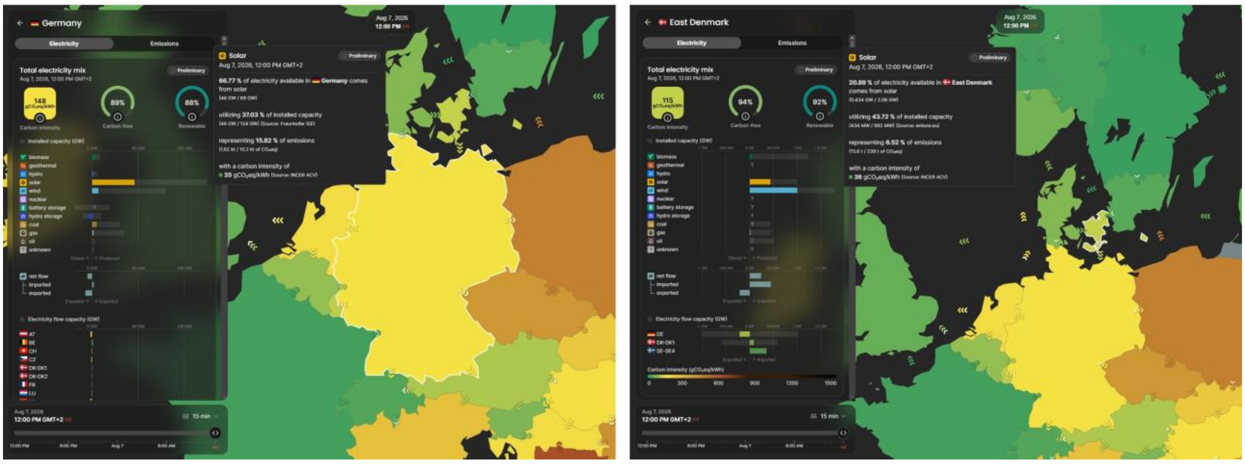

Figure S.4: *Germany and East Denmark at 12:00 GMT+2 on 7 August. The difference narrows to 6.69 percentage points, providing an internal counterexample to a fixed country ranking.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

GERMANY · 9 AUG · 12:00 GMT+2 · 41.96%

EAST DENMARK · 9 AUG · 12:00 GMT+2 · 93.84%

Figure S.5: *Germany and East Denmark at 12:00 GMT+2 on 9 August. The reported utilization gap is 51.88 percentage points.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

GERMANY · 10 AUG · 12:00 GMT+2 · 36.96%　　EAST DENMARK · 10 AUG · 12:00 GMT+2 · 27.11%

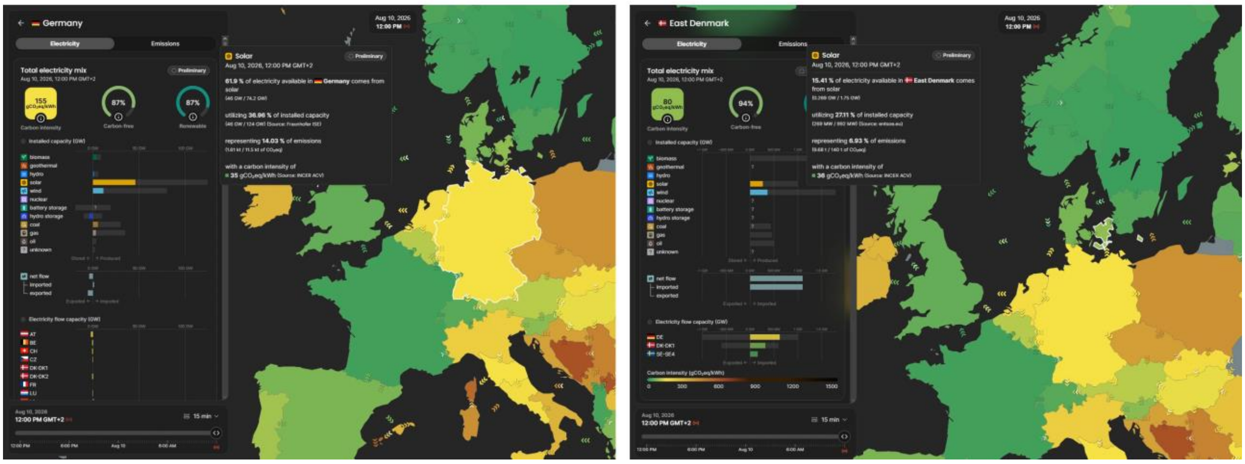

Figure S.6: *Germany and East Denmark at 12:00 GMT+2 on 10 August. East Denmark is 9.85 percentage points below Germany, demonstrating a reversal.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

GERMANY · 12 AUG · 13:30 GMT+2 · 46.60%　　POLAND · 12 AUG · 13:30 GMT+2 · 54.92%

Figure S.7: *Germany and Poland at 13:30 GMT+2 on 12 August. Poland is 8.32 percentage points higher.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

GERMANY · 13 AUG · 12:00 GMT+2 · 45.10%   EAST DENMARK · 13 AUG · 12:00 GMT+2 · 85.70%

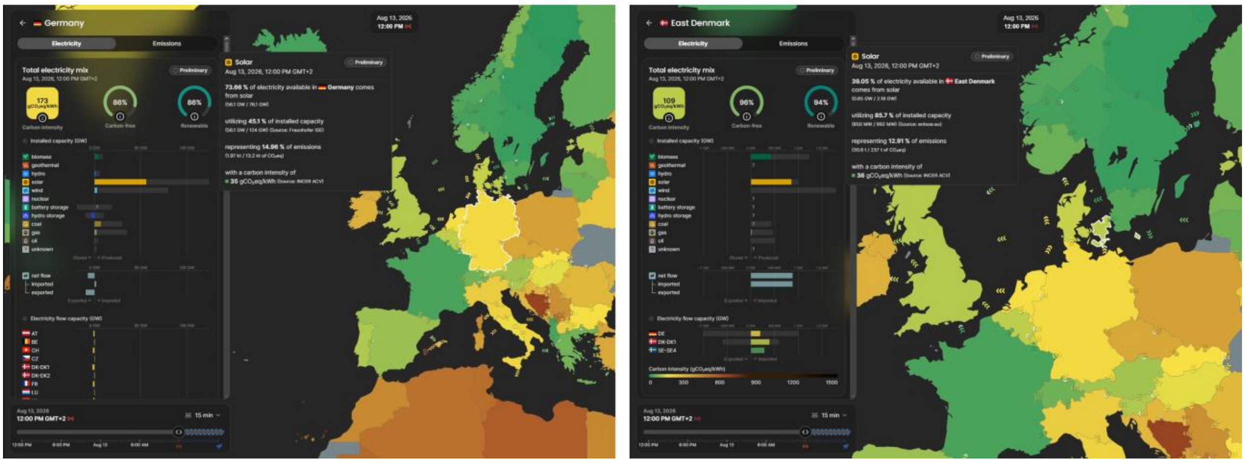

Figure S.8: *Germany and East Denmark at 12:00 GMT+2 on 13 August. East Denmark is 40.60 percentage points higher.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

GERMANY · 13 AUG · 12:00 GMT+2 · 45.10%   POLAND · 13 AUG · 12:00 GMT+2 · 65.10%

Figure S.9: *Germany and Poland at 12:00 GMT+2 on 13 August. Poland is 20.00 percentage points higher.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

S.6 Spain and Portugal

The Iberian observations show a persistent directional difference within the retained set: Portugal's reported output-to-listed-capacity ratio exceeds Spain's on every selected date. The magnitude varies from 27.80 to 57.11 percentage points. Values above 100% indicate a mismatch between reported output and the platform's installed-capacity denominator or the underlying AC/DC and reporting boundary; they are not evidence that a physical plant exceeded its own certified limit.

SPAIN · 5 AUG · 12:15 GMT+2 · 46.04%

PORTUGAL · 5 AUG · 12:15 GMT+2 · 95.83%

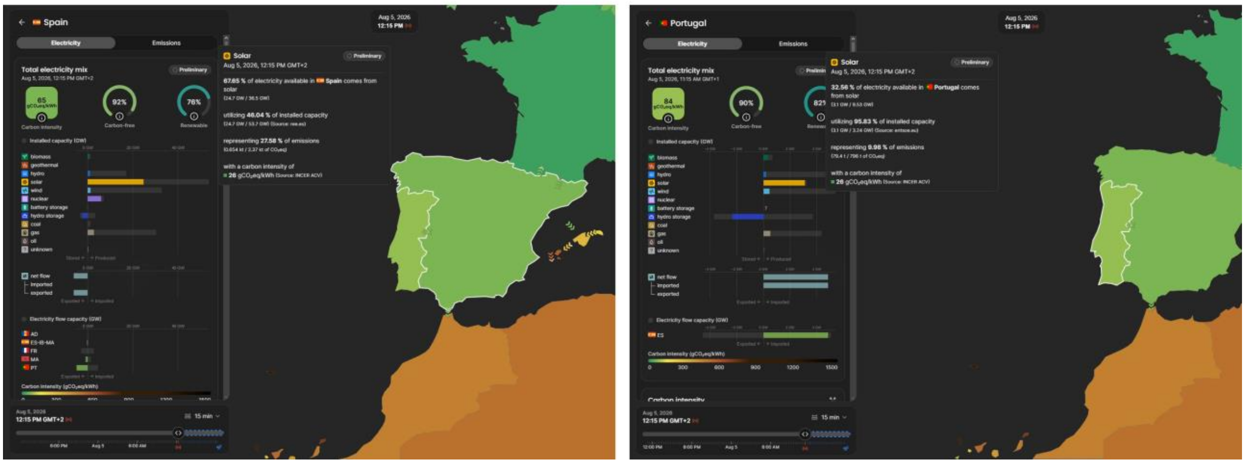

Figure S.10: Spain and Portugal at 12:15 GMT+2 on 5 August. Portugal is 49.79 percentage points higher.

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

SPAIN · 6 AUG · 12:45 GMT+2 · 49.43%

PORTUGAL · 6 AUG · 12:45 GMT+2 · 95.76%

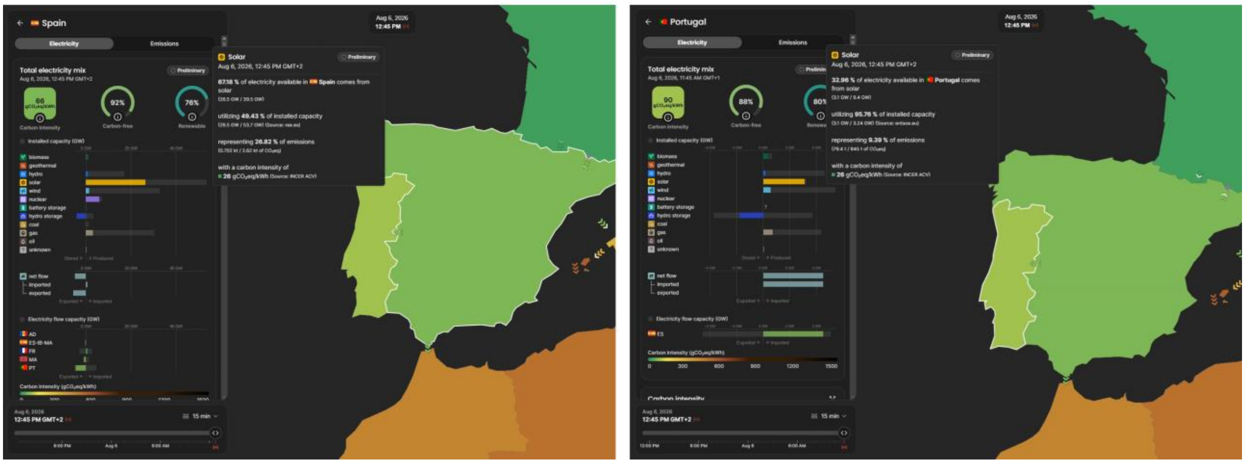

Figure S.11: Spain and Portugal at 12:45 GMT+2 on 6 August. Portugal is 46.33 percentage points higher.

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

SPAIN · 9 AUG · 12:00 GMT+2 · 42.45% | **PORTUGAL · 9 AUG · 12:00 GMT+2 · 70.25%**

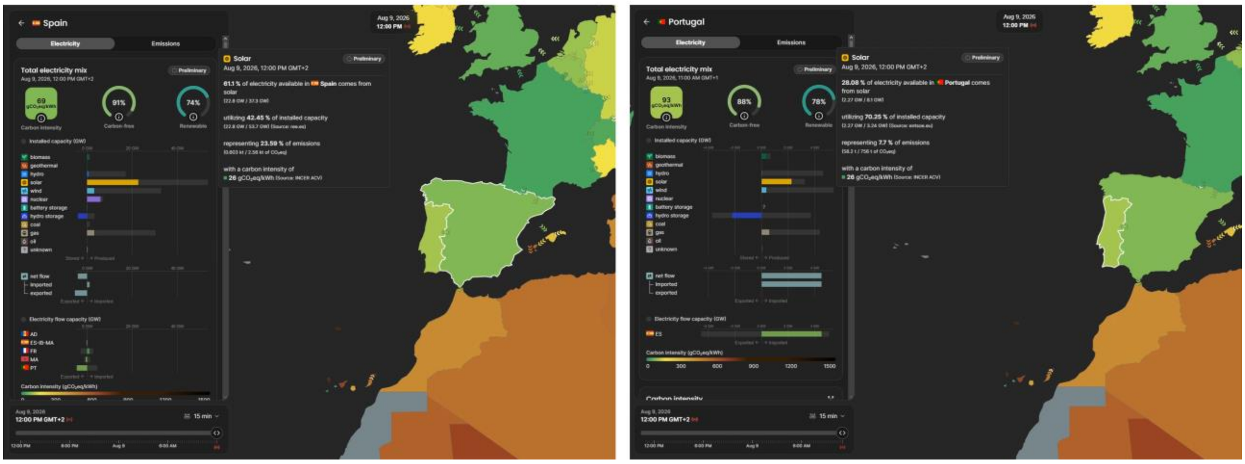

Figure S.12: *Spain and Portugal at 12:00 GMT+2 on 9 August. Portugal is 27.80 percentage points higher.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

SPAIN · 10 AUG · 12:00 GMT+2 · 48.00% | **PORTUGAL · 10 AUG · DISPLAY 12:00 GMT+2 · 100.74%**

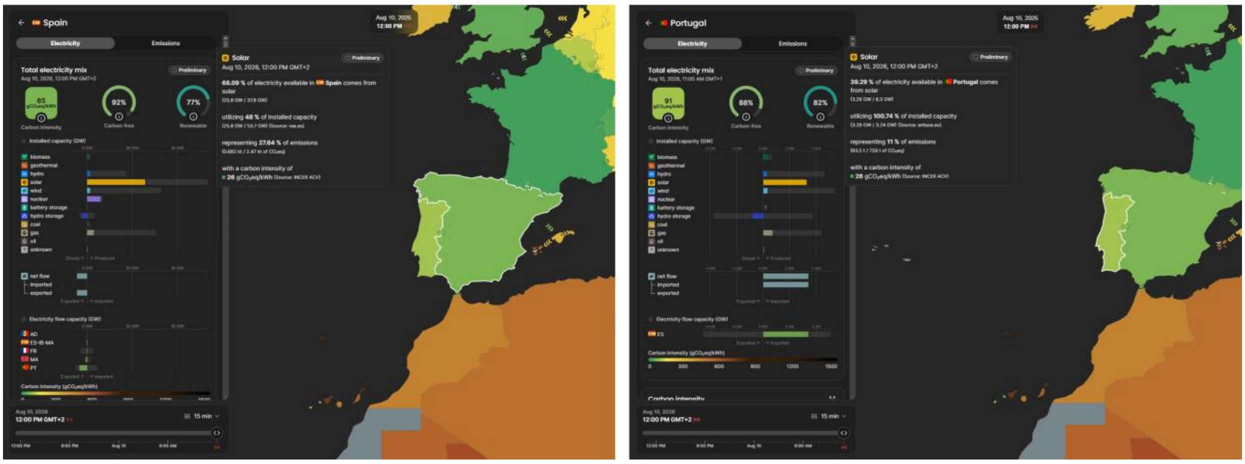

Figure S.13: *Spain and Portugal on the 12:00 GMT+2 display of 10 August. Portugal is 52.74 percentage points higher; the Portugal popup contains an internal time-label inconsistency.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

SPAIN · 11 AUG · 13:30 GMT+2 · 53.06%

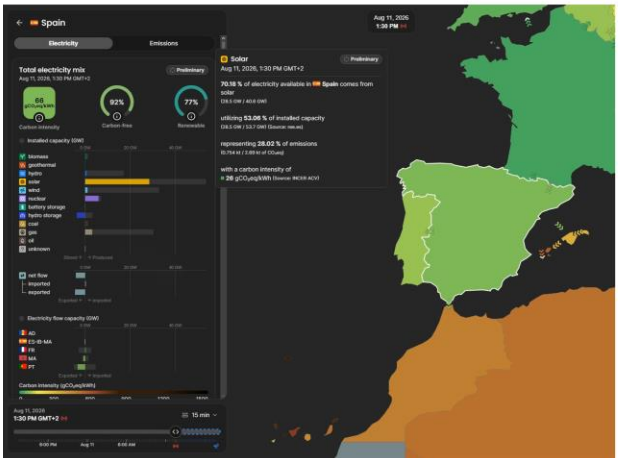

PORTUGAL · 11 AUG · 13:30 GMT+2 · 107.23%

Figure S.14: *Spain and Portugal at 13:30 GMT+2 on 11 August. Portugal is 54.17 percentage points higher.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

SPAIN · 13 AUG · 12:00 GMT+2 · 44.37%

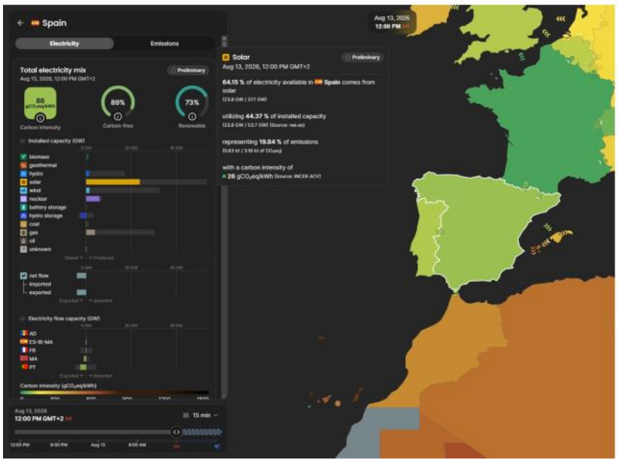

PORTUGAL · 13 AUG · 12:00 GMT+2 · 101.48%

Figure S.15: *Spain and Portugal at 12:00 GMT+2 on 13 August. Portugal is 57.11 percentage points higher.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

S.7 The recovered 8 August daily observation

The 8 August screenshot is a daily aggregate rather than an instantaneous midday observation. Electricity Maps displays 0.465 TWh of solar generation, 1.16 TWh of electricity available and a derived 2.98 TWh production potential. The resulting 15.58% ratio must therefore remain separate from the instantaneous values in Sections S.5 and S.6. Its value is continuity: it closes the date gap in the German observation sequence without creating a false like-for-like comparison.

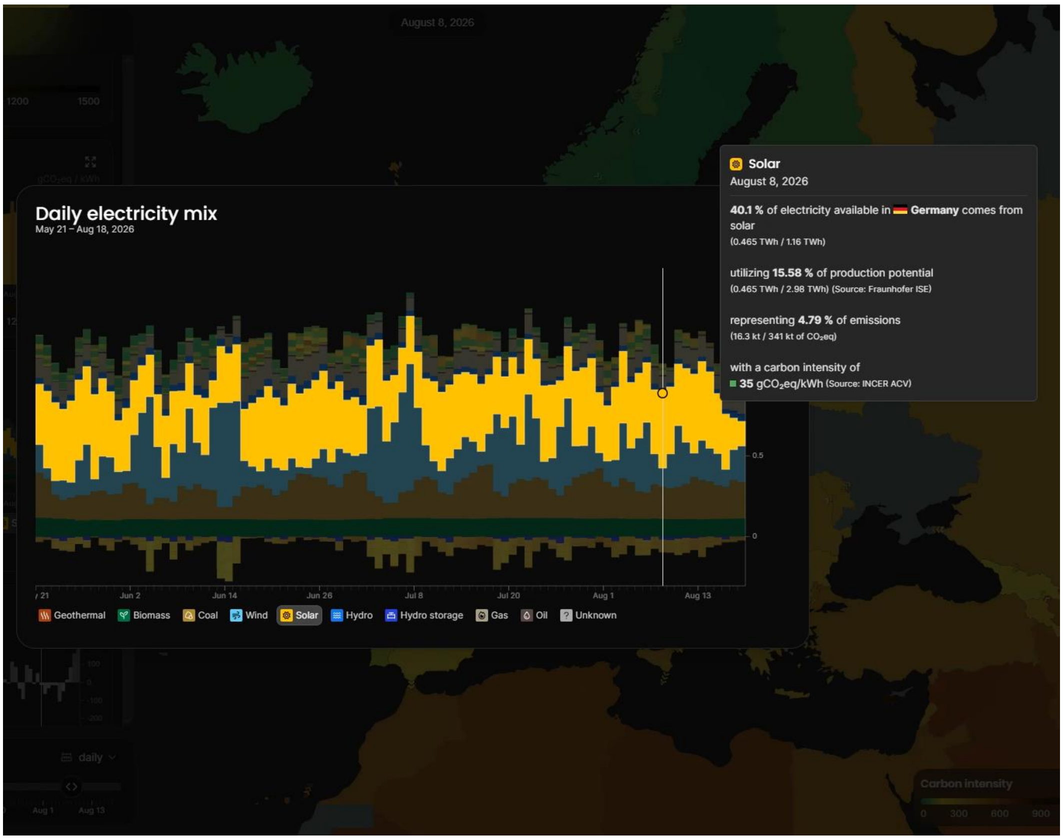

Figure S.16: *Germany daily solar aggregate for 8 August 2026: 0.465 TWh generation, 40.1% of displayed electricity available and 15.58% of the displayed daily production potential.*

Source: Electricity Maps screenshot retrieved by the author on 19 August 2026 for the historical date 8 August 2026.

S.8 Windy meteorological context

The Windy cloud-layer captures are retained as meteorological context for the European observations. They support the assessment of broad cloud fields and regional weather contrasts, but they are not pyranometer measurements and cannot establish irradiance at individual photovoltaic sites. Their evidentiary role is therefore contextual: they help test whether an obvious large-scale weather contrast is visible, while leaving plant-level solar resource and cloud-edge effects unresolved.

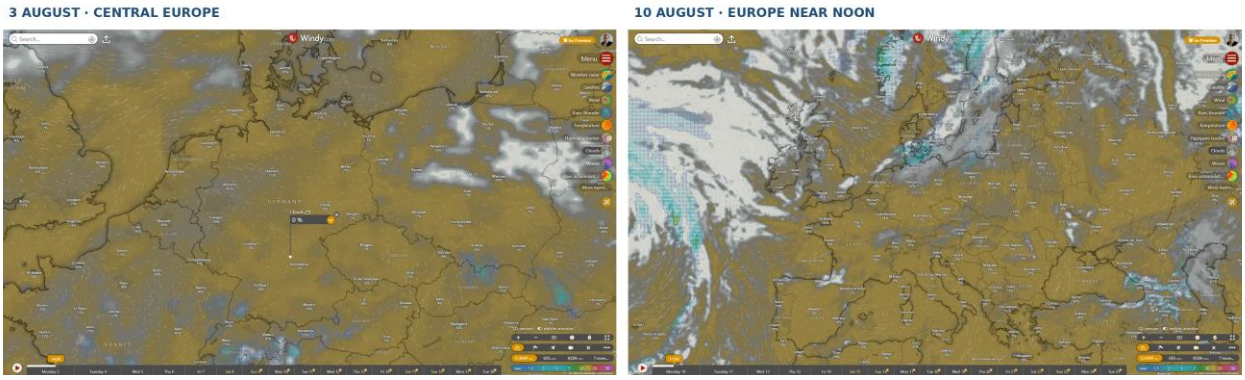

Figure S.17: *Windy cloud-layer observations for 3 and 10 August 2026. The images provide broad meteorological context for Central Europe; they do not measure site-level irradiance.*

Source: Windy screenshots supplied by the author. Displayed cloud percentages are model/interface values.

S.9 International control cases: ERCOT and California ISO

ERCOT and California ISO are not simultaneous weather controls for Europe. They are structural controls observed during the corresponding local solar-production window. The distinction is important: the European country pairs compare neighbouring systems at the same civil time, whereas the US observations test whether markedly different reported output-to-capacity ratios occur around local noon and, in ERCOT, well into the afternoon.

Date	Control area	Local time	Solar output	Listed capacity	Displayed ratio
4 Aug	ERCOT	2:30 PM CDT	31.4 GW	30.3 GW	103.65%
5 Aug	ERCOT	11:30 AM CDT	33.1 GW	30.3 GW	108.98%
11 Aug	ERCOT	4:45 PM CDT	29.2 GW	30.3 GW	96.25%
13 Aug	ERCOT	2:45 PM CDT	31.2 GW	30.3 GW	102.83%
7 Aug	California ISO	12:15 PM PDT	20.1 GW	24.9 GW	80.55%
13 Aug	California ISO	12:45 PM PDT	21.9 GW	24.9 GW	87.78%

Interpretation note. Values above 100% do not indicate a physical violation. They show that the reported output and the platform's listed installed-capacity denominator do not necessarily share an identical technical boundary, AC/DC definition or update date.

4 AUG · 2:30 PM CDT · 103.65% | 5 AUG · 11:30 AM CDT · 108.98%

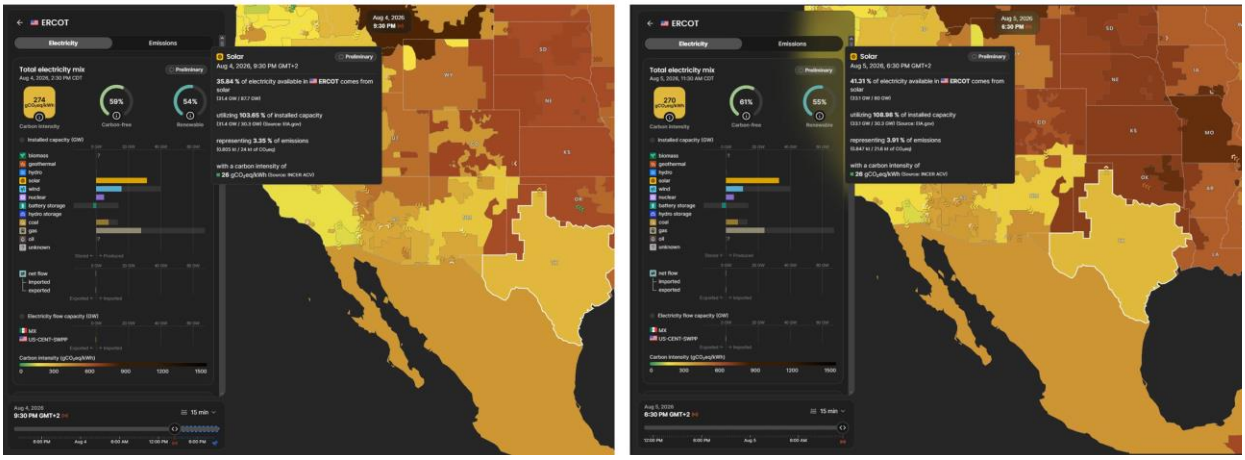

Figure S.18: *ERCOT around the local solar-production window on 4 and 5 August. Electricity Maps displayed ratios above 100% of its listed 30.3 GW capacity denominator.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

13 AUG · 2:45 PM CDT · 102.83% | 11 AUG · 4:45 PM CDT · 96.25%

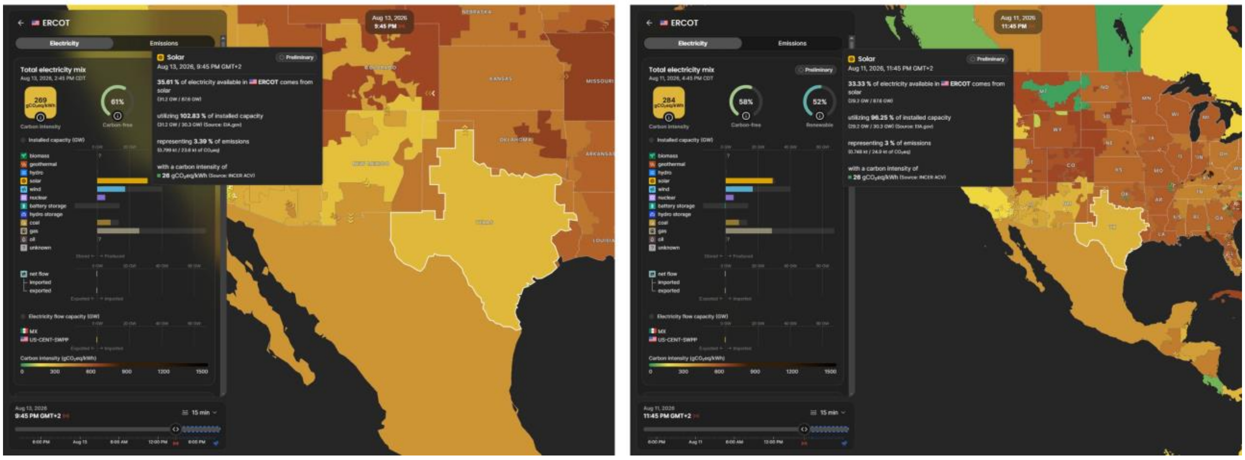

Figure S.19: *ERCOT afternoon observations: 102.83% at 2:45 PM CDT on 13 August and 96.25% at 4:45 PM CDT on 11 August. The high reported ratio was therefore not confined to solar noon.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

7 AUG · 12:15 PM PDT · 80.55% | 13 AUG · 12:45 PM PDT · 87.78%

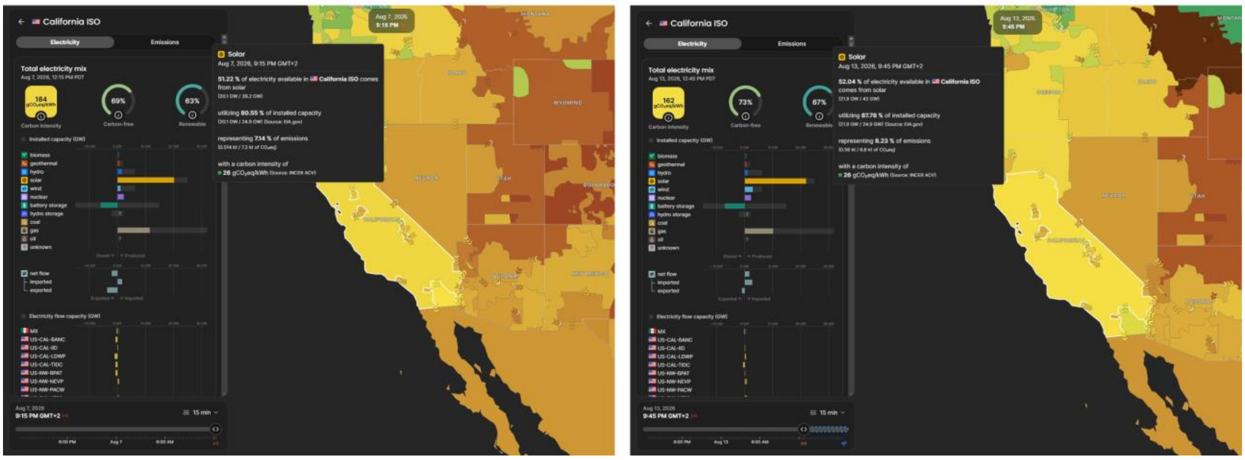

Figure S.20: *California ISO near local noon on 7 and 13 August. The platform displayed 80.55% and 87.78% respectively against its listed 24.9 GW capacity denominator.*

Source: Electricity Maps screenshots supplied by the author; preliminary platform data.

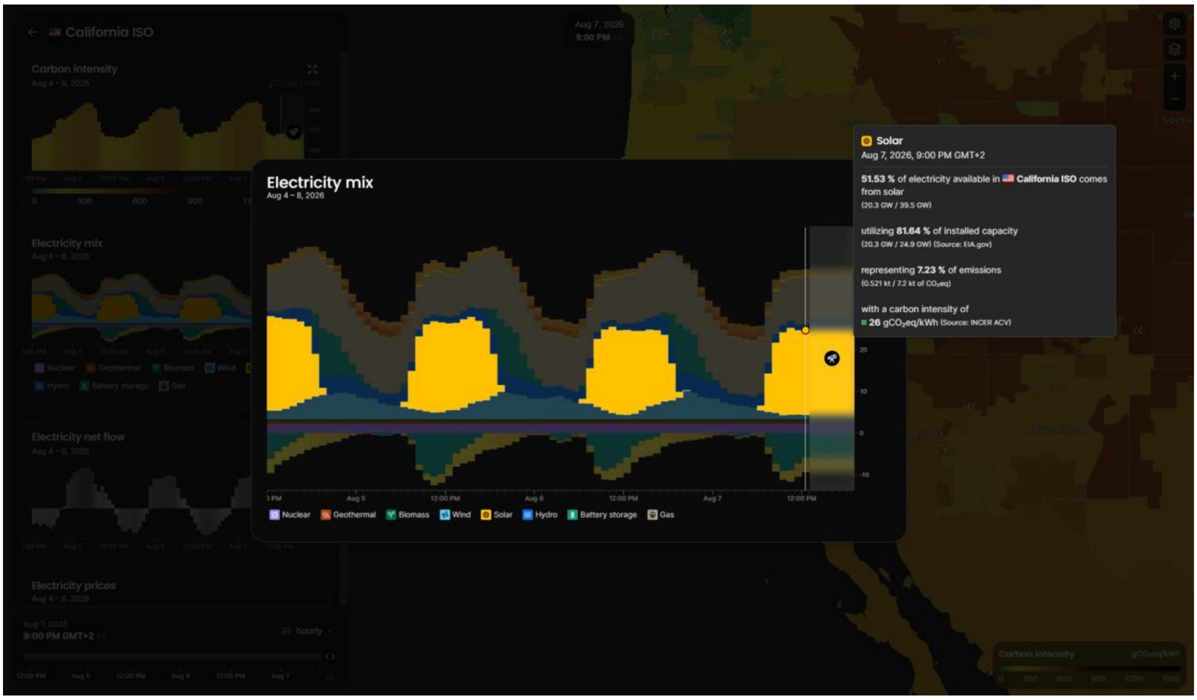

Figure S.21: *California ISO multi-day electricity-mix view for 4–8 August. It documents the recurring solar profile and battery-storage movement shown by the platform, but it is not a direct state-of-charge measurement and does not quantify curtailment.*

Source: Electricity Maps screenshot supplied by the author; preliminary platform data.

S.10 Candidate explanations and unresolved variables

The observations establish variation, not a unique cause. A complete explanation may require several variables at once: fleet orientation and tracking, DC/AC ratios, distributed self-consumption, plant availability, market dispatch, imports and exports, storage and pumped-hydro operation, network constraints, curtailment, data latency and differences in the installed-capacity denominator.

Candidate variable	Evidence needed	Status in this annex
Meteorology and irradiance	Satellite/cloud fields plus ground irradiance	Context only; not plant-level measurement
Fleet geometry and AC/DC design	Plant registry, orientation, tracker and inverter data	Not resolved
Storage and hydro buffering	Charging/pumping telemetry and reservoir dispatch	Visible in selected interfaces only
Cross-border flows and demand	Synchronized zonal flow and load data	Partially visible
Curtailment and redispatch	Operator instructions, baselines and plant meters	Outside the quantitative claim of this annex
Platform denominator and data age	Country-specific metadata and revision history	Material, especially where values exceed 100%

S.11 Relevance to the Expanded Gigawatt Fallacy

The empirical relevance to the NRG4NOW White Book is limited but direct. Installed capacity in gigawatts does not by itself determine the energy admitted into, absorbed by or made available from an electricity system at a particular time. Neighbouring systems with connected markets and the same civil time can display materially different output-to-capacity ratios. The observation therefore supports the White Book's distinction between nameplate capacity, delivered energy, temporal availability and usable system access.

The annex does not establish that any one country's lower ratio represents avoidable loss, nor does it calculate a national volume of recoverable energy. It identifies a measurable divergence that warrants examination with primary operational data and that justifies architectures designed around availability rather than nameplate capacity alone.

S.12 Source register

Electricity Maps. Time-stamped electricity-mix, installed-capacity and source-detail interface screenshots for 3–13 August 2026. Data labels include Fraunhofer ISE, ENTSO-E, REE and EIA depending on the selected zone. Screenshots supplied by the author; most interface values marked preliminary.
Windy.com. Cloud and weather-layer screenshots for selected observation dates. Used only as broad meteorological context, not as site-level irradiance measurements.
Author's archive: 3-13.zip, 65 files, supplemented by DE_08.PNG retrieved for 8 August 2026. Review set: 66 files.

S.13 Editorial conclusion

Across the retained European observations, large differences appear repeatedly between Germany, East Denmark and Poland and between Spain and Portugal. The differences are not constant, and the 10 August Germany–East Denmark reversal demonstrates why a fixed national ranking would be unjustified. Windy adds broad meteorological context without replacing irradiance measurements. The ERCOT controls extend the finding beyond local noon: Electricity Maps displayed 96.25% at 4:45 PM CDT on 11 August and 102.83% at 2:45 PM CDT on 13 August. California ISO supplies a further local-noon and storage-cycle control. The defensible conclusion remains narrow: timestamp, geographic proximity and nameplate capacity alone are insufficient to explain the reported output and system response. The screenshots identify a material divergence for further examination; they do not by themselves quantify curtailment.