



rehlko

A Structured Transition Model for AI Data Center Power

Designing Reliable Infrastructure Today Without Locking Out Tomorrow's Decarbonization



Disclaimer

This paper is intended for informational and strategic planning purposes only. It reflects illustrative system concepts and modeled outcomes based on representative assumptions. It does not constitute engineering advice, a performance guarantee, or a commitment regarding any specific project or outcome. Certain statements in this paper describe potential performance, including emissions outcomes, system efficiency, and operational characteristics. Actual results will vary based on site conditions, equipment configuration, fuel supply, regulatory requirements, and operational practices. Lifecycle carbon values and comparative results are illustrative and should not be used for regulatory reporting, project finance, or investment decisions without project-specific analysis and validation.



Executive Summary

AI-driven data centers are expanding faster than the power systems built to support them. Interconnection delays, volatile energy markets, and rising decarbonization expectations are forcing developers to make long-term infrastructure decisions under significant uncertainty. At the same time, AI workloads require exceptional reliability and rapid deployment, creating a fundamental tension between speed, resilience, and sustainability.

This paper introduces a Structured Transition Model: a lifecycle-oriented approach to data center power that delivers reliable capacity today while preserving the flexibility to hybridize, decarbonize, and integrate with future grid conditions. Importantly, the Structured Transition Model treats speed to power as a sustainability enabler. It allows efficient compute capacity to come online immediately using proven, fuel-flexible platforms that meet today's emissions requirements while remaining fully compatible with future decarbonization pathways. Rather than optimizing for a single end state, the model defines key performance characteristics: modularity, fuel flexibility, part-load efficiency, and system interoperability, that enable infrastructure to evolve over time.

The model is structured across three phases: establishing a reliable initial speed to power foundation; integrating renewable supply, storage, and thermal recovery; and transitioning fuels while increasing grid interaction. Projects may enter at any phase depending on site conditions and strategy, provided reliability and future optionality are maintained.

Lifecycle performance is central to this approach. Systems designed for staged evolution can outperform static designs by reducing political risk, stranded asset risk, preserving capital value, and enabling progressive decarbonization over a 15–20 year horizon.

Hybridization plays a critical role, improving efficiency, lowering carbon intensity, and enabling participation in flexibility and capacity markets. As systems evolve, dispatchable generation shifts from continuous operation toward a balancing and resilience role within increasingly dynamic power systems.

The model emphasizes adaptability over prescription. Rather than predicting future policy or market conditions, it prioritizes infrastructure that remains viable across a range of scenarios.

For developer and investors, this enables faster deployment with reduced long-term risk.

AI data centers are not static assets but evolving energy systems. The most resilient strategy is structured evolution: reliability today, flexibility tomorrow, and decarbonization over time.

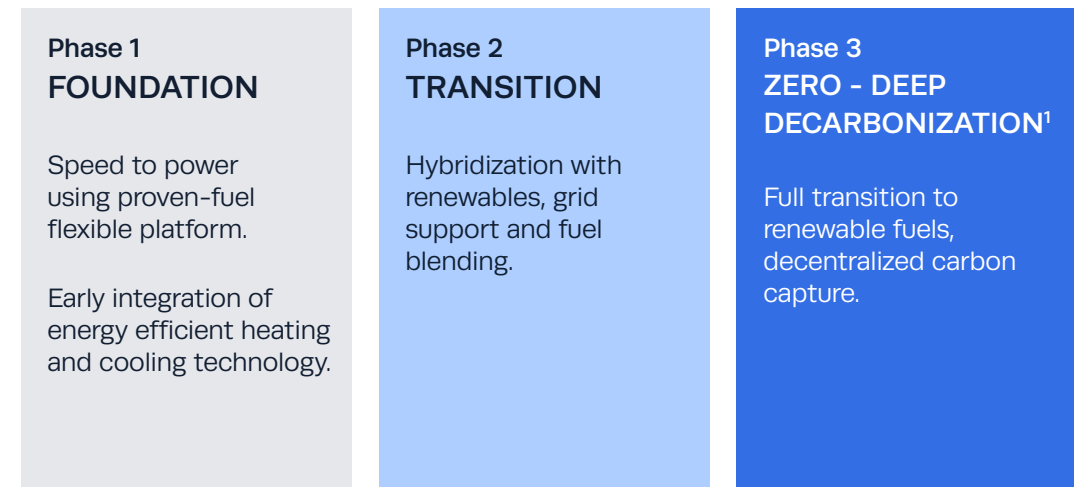
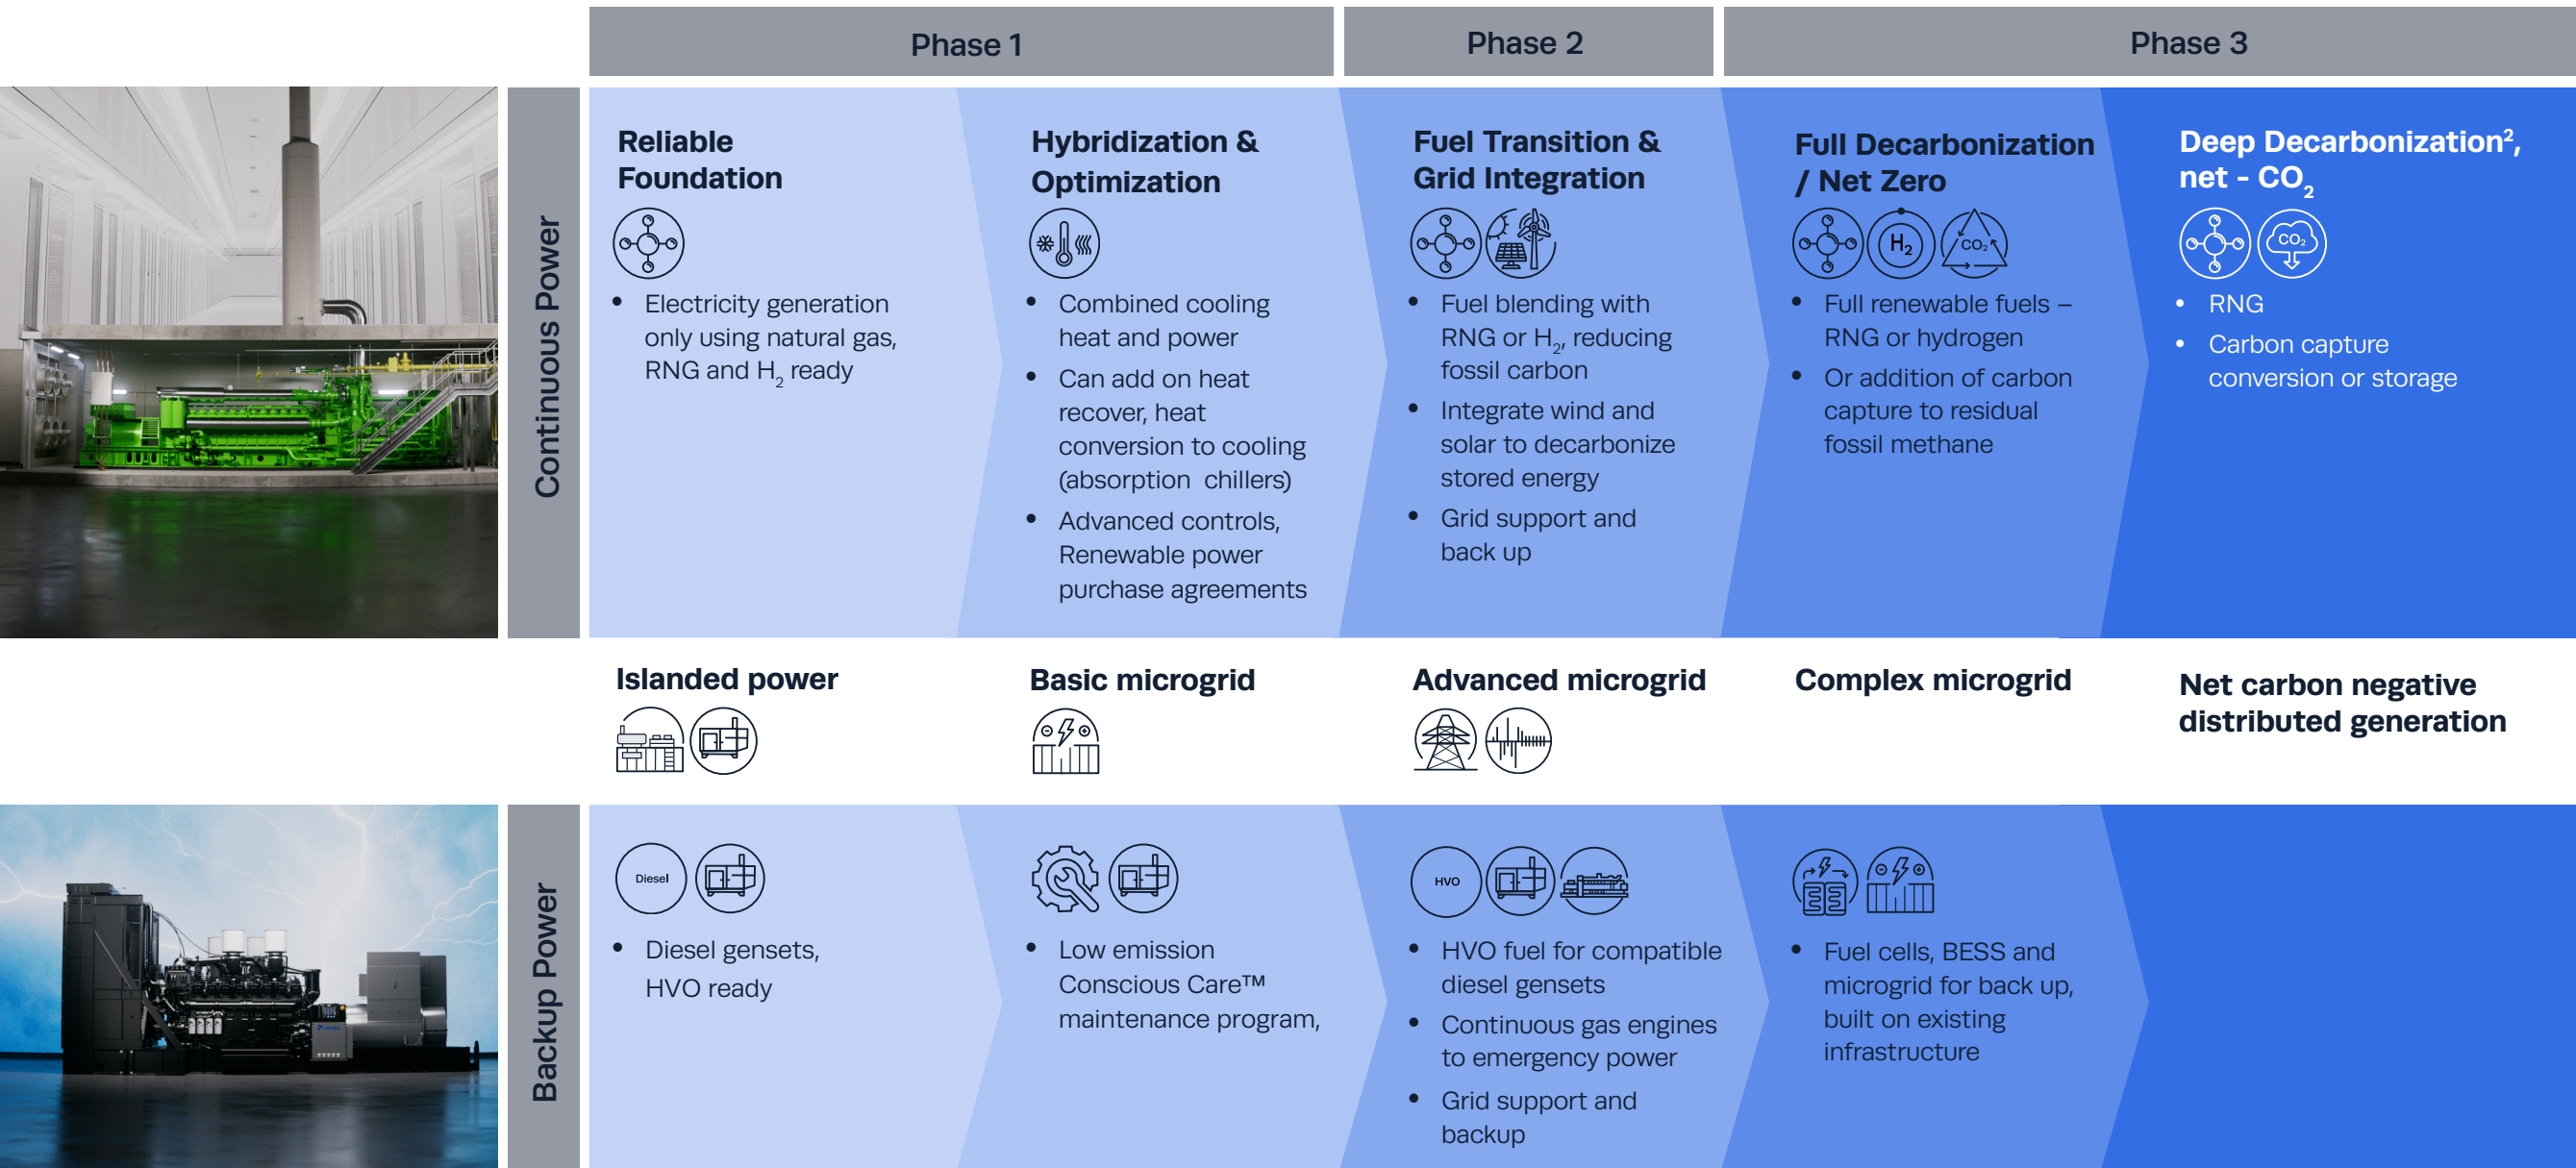


Figure 1: The Structured Transition Model at a Glance

¹ Depending on feedstock, system configuration, and lifecycle accounting methodology, certain pathways may achieve very low or potentially net-negative lifecycle emissions; such outcomes are highly dependent on project-specific assumptions and require independent validation.

Suite of Sustainability Solutions for Data Centers



² Depending on feedstock, system configuration, and lifecycle accounting methodology, certain pathways may achieve very low or potentially net-negative lifecycle emissions; such outcomes are highly dependent on project-specific assumptions and require independent validation.

Figure 2: The Structured Transition Model Sub-Phases and Key Components

Contents

Context: Why Traditional Power Planning Is Failing	6
AI Load Profiles & Power System Design	8
Principle: Speed to power, future flexibility	10
Reliability and the Structured Transition Model	13
Phase 1a: Deployment and Reliability Foundation (Years 0–5)	15
Phases 1a-2: Hybridization and Optimization (Years 3–10)	16
Phases 2-3: Fuel Transition and Grid Integration (Years 7–20)	20
Lifecycle Carbon Impact of Alternative Power Architectures	24
Political Risk & Infrastructure Longevity	30
Conclusion	33



Context:

Why traditional power planning is failing

AI workloads are changing the assumptions on which power systems have traditionally been planned. Power densities are rising sharply as compute intensity increases, while load profiles are becoming more volatile and less predictable. Interconnection timelines often stretch beyond construction schedules, creating a persistent mismatch between when capacity is needed and when firm grid access can be secured. At the same time, carbon expectations are shifting from a narrow focus on operational emissions to a broader consideration of lifecycle impact, including how infrastructure is built, fueled, dispatched and refreshed over time.

This aligns with IEA analysis³ showing that AI-driven data centers are set to become one of the fastest-growing sources of electricity demand globally, with grid capacity and infrastructure, not compute - emerging as the primary constraint on scale.

In this environment, waiting for a perfect zero carbon solution frequently results in non deployment, deferring both economic and environmental value along with digital capacity. Developers and investors are too often confronted with a false binary: move fast with fossil heavy, energy inefficient options and accept elevated emissions, or pursue renewables only strategies that may delay projects and leave reliability uncertain. Neither approach reflects how complex infrastructure systems evolve in practice. The more robust alternative is to plan intentionally for staged development that preserves future optionality while meeting near term demand and reliability targets.

To frame the practical implications of current power strategies, it is helpful to compare the dominant approaches available to developers today. Each presents a different balance of deployment speed, carbon performance, capital risk, and system flexibility over the asset lifecycle.

³ International Energy Agency (2025) – Energy and AI

Comparative Landscape of Data Center Power Strategies

	Grid-Only (PPA / Utility Supply)	Diesel Backup Model	Structured Transition Model
Time to Power	Long (3–7+ yrs) Dependent on interconnection and transmission upgrades	Short (6–18 m) Fast deployment of backup capacity, but limited prime capability	Short to Medium (9–24 m) On-site dispatchable generation enables early operation, scalable with demand.
Lifecycle CO₂ Intensity	Low to Moderate (declining over time) Dependent on grid mix and power purchase agreement (PPA) structure	High (peaking-only use, high emissions when dispatched)	Moderate → Low (declining over lifecycle) Improves through hybridization, efficiency gains, and fuel transition.
Capex at Risk (Stranding Exposure)	Low initial capex, but high dependency risk Exposure to grid delays and future pricing	Moderate to High Backup assets may become underutilized or non-compliant under tighter regulations	Low to Moderate Modular, fuel-flexible assets retain value across phases; limited need for replacement
Flexibility Revenue Potential	Limited Typically passive consumer with minimal market participation	Very Limited Assets rarely operate and are not optimized for market participation	High Access to capacity markets, ancillary services, demand response, and energy arbitrage through hybrid systems

AI Load Profiles & Power System Design

AI data centers are changing not just the scale of demand, but the nature of electrical load itself. Traditional data center design has been based on relatively stable, predictable utilisation patterns, with gradual ramp-up and limited short-term variability. AI-driven workloads introduce a different profile: high power density, rapid scaling, and increasingly dynamic behaviour at both cluster and facility level.

Training workloads can create step changes in demand as compute clusters are brought online, while inference workloads introduce variability tied to user activity and model deployment cycles. As clusters are reconfigured or expanded, power demand can shift materially over short timeframes. The result is a load profile that is less linear, less predictable, and more sensitive to infrastructure constraints than in previous generations of data center development.

This has two direct implications for power system design.


First, systems must be capable of accommodating non-linear load growth. Capacity is often required ahead of full utilization, and infrastructure must be able to scale in modular increments without

compromising efficiency or reliability. Large, indivisible generation assets or architectures that depend on long lead-time grid upgrades are poorly aligned with this requirement. By contrast, modular dispatchable generation; deployed in parallel with IT load growth, allows capacity to be added in step with demand while maintaining operational flexibility.

Second, power systems must respond effectively to intra-day variability and rapid ramping. AI workloads can exhibit sharper changes in demand than

traditional enterprise or cloud applications, placing greater emphasis on part-load efficiency, fast start capability, and dynamic response. Technologies that are optimised for steady-state operation at or near full load may experience efficiency penalties, increased wear, or operational constraints under these conditions.

Figure 3: An idealised load profile and battery operation over 60 minutes (balanced at 5MWh) - Lenoff (2025)



AI Load Profiles & Power System Design, cont.

The Structured Transition Model is inherently aligned with these characteristics. In its early phases, distributed generation provides a stable and controllable supply capable of meeting both base and variable load. As systems evolve, battery storage and advanced control platforms are introduced to manage short-duration fluctuations, smooth load transitions, and optimize dispatch across multiple energy sources. Dispatchable engines transition from continuous operation toward a balancing role, complementing variable renewable supply and responding to changes in demand.

Crucially, this approach treats variability not as an exception to be managed, but as a defining feature of the system. By combining modular generation, storage, and intelligent control, the Structured Transition Model enables power infrastructure to scale and respond in line with the evolving behaviour of AI workloads, while maintaining the reliability envelope required for continuous operation.



Principle:

Speed to power, *future flexibility*

The Structured Transition Model rests on a single engineering principle: infrastructure should be designed as a flexible transition pathway rather than as a fixed endpoint. The Structured Transition Model is a conceptual framework intended to support strategic planning and does not replace detailed engineering, commercial, or regulatory analysis for specific projects.

In practical terms, this means starting with proven dispatchable reliability; enabling modular hybridization over time; preserving the fuel flexibility of deployed assets; and facilitating grid integration as and when it becomes available.

Crucially, the model is not defined by any single technology class, but by the ability of selected assets to meet key system requirements: including fast response, efficient part-load operation, compatibility with multiple fuel pathways, and integration within a dynamically controlled microgrid environment.

The objective is not perfection on day one but a credible decarbonization trajectory over the life of the asset. One that avoids stranded infrastructure, maintains a high reliability envelope and supports continuous, low friction improvement.

While illustrated as a phased progression, real-world deployments are expected to be non-linear: assets may shift between operating modes, delay transitions, or accelerate hybridisation depending on market conditions, fuel availability, and policy signals.

Principle: Speed to power, future flexibility

Aligning Transition Pathways with Industry Power Maturity

The Structured Transition Model outlined in this paper focuses specifically on the evolution of energy systems that support AI data centers. While intentionally practical and engineering led, it aligns closely with the broader industry understanding of sustainability maturity as defined by the iMasons Climate Accord (iCA) Maturity Model.²

The iCA framework establishes a portfolio level view of decarbonization maturity across Power, Materials, and Equipment, with its Power Maturity Model articulating progressively higher levels of performance; from annual renewable matching through to hourly matching, carbon free locality, and active grid benefit. Importantly, the iCA model is deliberately non prescriptive, defining what higher maturity looks like without dictating how those outcomes must be achieved.

The Structured Transition Model can be viewed as a complementary execution framework for the "Power" dimension of this maturity model. As data center operators move toward higher maturity levels, innovation shifts away from individual technologies and toward integrated systems, combining dispatchable generation, storage, controls, fuels, and grid interaction. It is at this

system level that the Structured Transition Model provides clarity: translating maturity objectives into coherent, deployable power architectures that preserve reliability while enabling staged decarbonization.

In this sense, the model supports accelerated progress through the upper tiers of power maturity by coupling industry defined outcomes with transition ready engineering design.

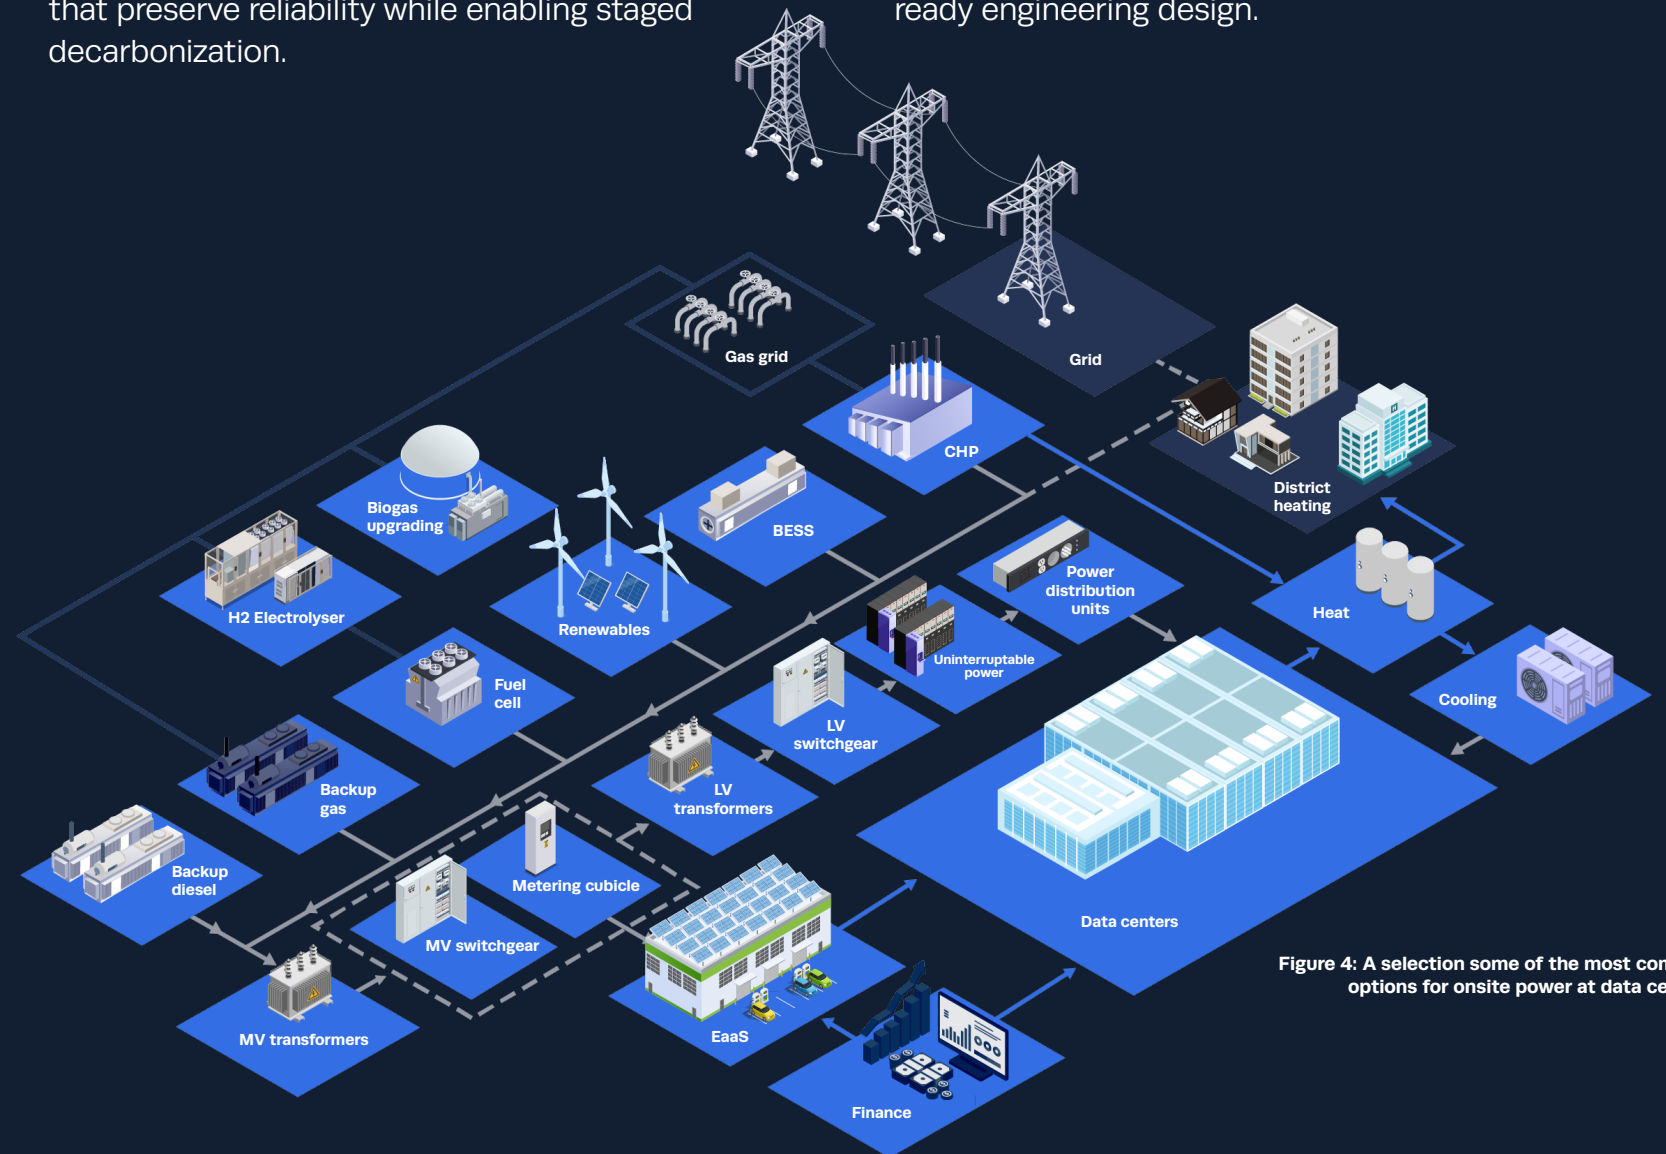


Figure 4: A selection some of the most common options for onsite power at data centers

⁴ Masons (2025) Climate Accord Maturity Model; A Guide to Emission Reduction in Digital Infrastructure

Comparative Landscape of Data Center Power Strategies

Site Condition	Typical Characteristic	Recommended Entry Phase	Rationale
Grid-constrained hyperscale	Long interconnection delays, limited firm capacity, urgent deployment timelines	Phase 1 or 2	Immediate on-site dispatchable generation enables rapid deployment, with hybridization introduced as grid and market conditions evolve
Grid-Rich / Low-Carbon Regions	Strong grid access, high renewable penetration, PPAs readily available	Phase 2 or 3	Hybrid or low-carbon systems can be deployed from day one, provided reliability is maintained
Remote / Islanded / Emerging Markets	Weak or unreliable grid, limited infrastructure, high energy cost volatility	Phase 1 → Extended Phase 2	On-site generation forms the backbone of supply, with hybridization improving efficiency and reducing fuel dependency over time
Industrial Co-Location / Heat Demand	Significant thermal loads, opportunity for CHP/CCHP integration	Phase 2	Early integration of thermal recovery maximizes system efficiency and improves lifecycle carbon performance

Reliability and the Structured Transition Model

Reliability remains the non-negotiable requirement for AI data center power. While efficiency and carbon performance continue to rise in importance, they are secondary to maintaining uninterrupted operation. The Structured Transition Model is therefore grounded in a simple premise: decarbonization must not come at the expense of uptime.

In traditional grid-reliant architectures, reliability is largely externalized. Availability depends on transmission infrastructure, generation adequacy, and grid stability—factors outside the direct control of the operator. Diesel backup systems provide resilience but are typically designed for infrequent operation and limited duration, rather than sustained or flexible use.

The Structured Transition Model takes a different approach by internalizing reliability within the site boundary. Dispatchable generation is deployed as part of the primary architecture, not simply as contingency. This allows data centers to operate independently of grid constraints when required, while maintaining full synchronization capability when grid supply is available.

Reliability is achieved through modular, redundant system design. Rather than relying on a small number of large assets, capacity is distributed across multiple generation units, typically configured with N+1 or higher redundancy. This architecture allows individual units to be taken offline for maintenance without compromising overall availability and reduces the impact of single-point failures. When combined with short-duration storage and advanced control systems, the result is a highly resilient, fault-tolerant power system capable of meeting stringent uptime requirements.

From a performance perspective, this approach supports availability levels aligned with Tier III and Tier IV data center standards. Well-designed multi-engine installations, supported by appropriate redundancy and maintenance strategies, can achieve availability in the range of four to five nines (99.99%–99.999%), depending on configuration and operational discipline.

Importantly, reliability within the Structured Transition Model is not static. As systems evolve through hybridization and grid integration, the role of dispatchable generation shifts; from continuous prime power toward balancing, resilience, and

contingency support, without reducing the overall reliability envelope. Battery storage improves transient response and ride-through capability, while grid connection, when available, provides an additional layer of supply diversity.

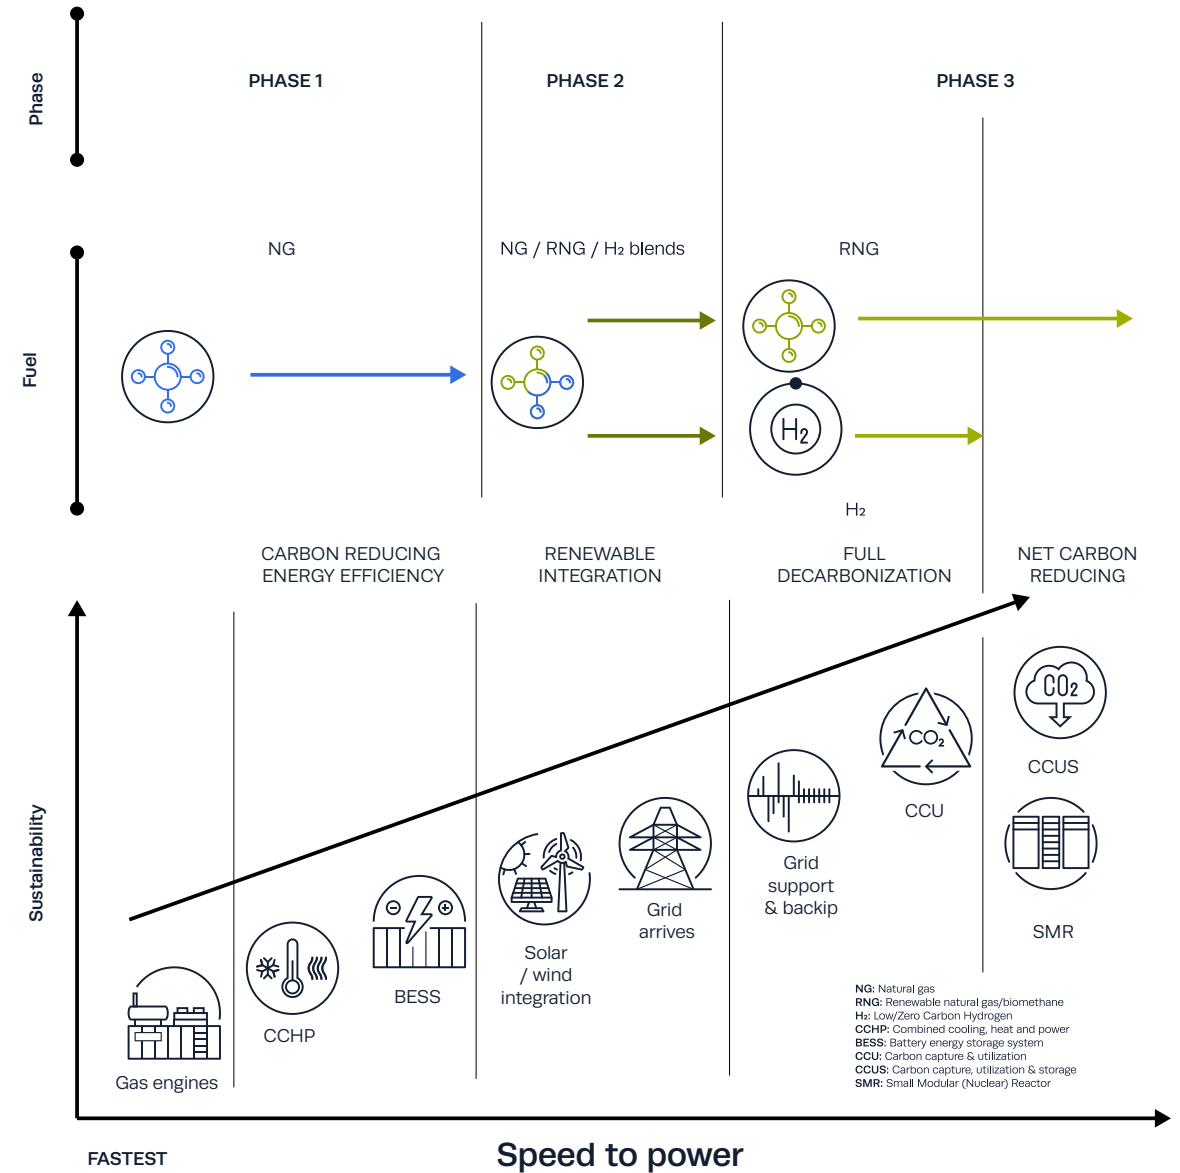
This layered approach to reliability, combining on-site generation, storage, and grid interaction, creates a system that is inherently more robust than any single-source configuration. It also ensures that as decarbonization progresses, reliability is preserved rather than traded away.

In this context, reliability is not simply designed in at commissioning; it is maintained and reinforced over the life of the asset.

Three Phases: The Structured Transition pathway

The pathway is presented in three phases for clarity. Projects do not need to begin at Phase 1. Developers can start at whichever phase matches site conditions and strategy, provided the design remains compatible with upstream and downstream evolution.

Figure 5: Technical options for a gas engine based structured transition pathway. The same initial prime power assets can be used for the full spectrum of decarbonization.



Three Phases: The Structured Transition Pathway

Phase 1a:

Deployment and Reliability Foundation (Years 0–5)

Phase 1a is not a stopgap. It is a deliberate design choice: deploying existing, bankable platforms that deliver immediate reliability and autonomy while preserving full compatibility with hybridization, renewable fuels, and grid integration. Speed to market is achieved without locking in emissions."

Three design principles underpin this foundation. First, generation assets should be selected for their ability to transition to lower carbon fuels in future, ensuring that today's investments remain relevant tomorrow. Second, bespoke architectures that inhibit later integration of electrical and thermal technologies should be avoided; the system should not be engineered into a corner. Third, the electrical topology should be prepared from the outset for multi source inputs, making it straightforward to incorporate additional generation, storage or grid services in later phases. Success in this first stage is measured by rapid time to operation, high availability and resilience, and a scalable architecture that can be expanded or adapted with minimal rework. Some developments will commence directly in later phases; for example, sites with robust low carbon grid access and immediate power (PPA) or renewable gas purchase agreements (RGPA) options may bypass Phase 1 and enter at Phase 2 or Phase 3, provided reliability criteria are met.

Figure 6: Technical options structured transition pathway for back up generation. Significant decarbonization can be achieved with existing gensets, with prime assets being transitioned to back up power for deepest levels of decarbonization.

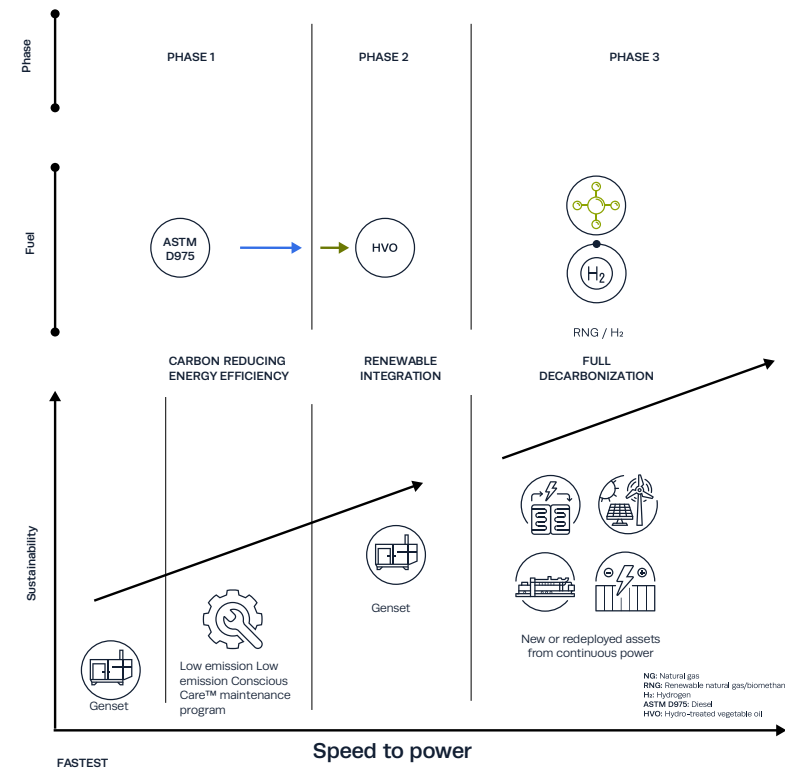
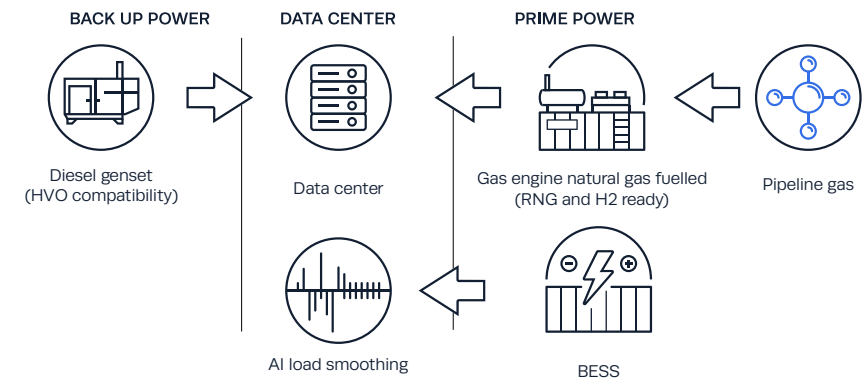


Figure 7: Deploying for rapid power - diesel and gas gensets able to be converted to renewable fuels, BESS option to support AI load smoothing.



Three Phases: The Structured Transition Pathway

Phases 1a-2:

Hybridization and Optimization (Years 3–10)

Once stable operations are established, the system can be hybridized to reduce emissions and costs while increasing flexibility, integrating renewable electricity, battery storage and, critically, combined cooling, heat and power (CCHP) architectures that treat power and cooling as a single, optimised energy system.

This system-level integration of power, cooling and heat recovery aligns with the Cogen World Coalition's 2025 analysis , which positions CHP and CCHP as enabling technologies for scalable, efficient and resilient data center development under AI-driven load growth and grid constraints.

At this stage, the economic and operational drivers of transition become increasingly important. Hybridization is typically triggered by a combination of factors including declining storage costs, increasing carbon pricing exposure, improved availability of renewable PPAs or renewable gas contracts, and the monetisation of flexibility through grid-service markets.

Renewable electricity supplies can be practically integrated through PPAs, adding renewable gas via RGPAs or on site production where practical, and integrating hydrotreated vegetable oil to de-fossilize backup generators.

Battery energy storage systems (BESS) with robust battery and energy management systems (BMS/EMS) are introduced to smooth intermittency and support peak shaving, while also enabling participation in capacity markets, demand response, and ancillary services, creating additional revenue streams that improve project economics.

Heat recovery and thermal storage from prime engines enable useful heat delivery to local users or district energy systems. Absorption chillers can convert otherwise unused thermal energy into cooling, improving overall efficiency. Absorption chillers are thermally driven cooling systems that use heat, rather than electricity, as their primary energy input, typically from waste heat, steam, hot water from CHP units. Using a refrigerant-absorbent pair, they convert thermal energy into chilled water for air conditioning or process cooling. Within a Structured Transition Model, absorption chillers enable electrification reduction, improve overall system efficiency through heat recovery, and support decarbonization pathways by valorizing surplus or low-grade heat that would otherwise be lost, making them particularly suited to data center applications.

Advanced control systems and microgrid orchestration manage these resources dynamically, and participation in grid services allows the site to capture value from flexibility markets.

⁵ Cogen World Coalition (2025) CHP – A Key Enabling Technology for Ensuring Data Center Development is Efficient, Cost-Effective and Resilient, https://cogenworld.org/wp-content/uploads/2025/12/Data_Center_Cogeneration_Opportunities_CWC_2025_12_04_1_.pdf

Three Phases: The Structured Transition Pathway

Operationally, dispatchable engines migrate from continuous duty towards a balancing and resilience role, providing capacity when renewable supply is limited or grid conditions deteriorate. Storage assets mitigate short term variability and improve economics by arbitraging tariffs and reducing fuel consumption. As the control layer matures, the system continuously optimises the energy mix and export or import decisions against reliability, cost and carbon objectives. Progress in this phase is reflected in falling carbon intensity per megawatt hour, improving Power Usage Effectiveness (PUE), reduced fuel consumption and increasing revenues from flexibility services. Many new builds can begin at Phase 2—for example, where PPAs, RGPAs and BESS can be contracted at practical cost—so long as the design retains forward and backward compatibility with adjacent phases.

Value Stacking Through Grid Interaction

As data center energy systems move beyond single-source supply models, interaction with the electricity grid becomes a key lever for both economic optimization and overall system value. Modern hybrid power architectures are no longer passive electricity consumers; they can operate as active grid participants, unlocking multiple, complementary revenue streams.

This capability—commonly referred to as value stacking—allows a single, integrated asset base to deliver resilience, decarbonization progress, and attractive financial returns in parallel.

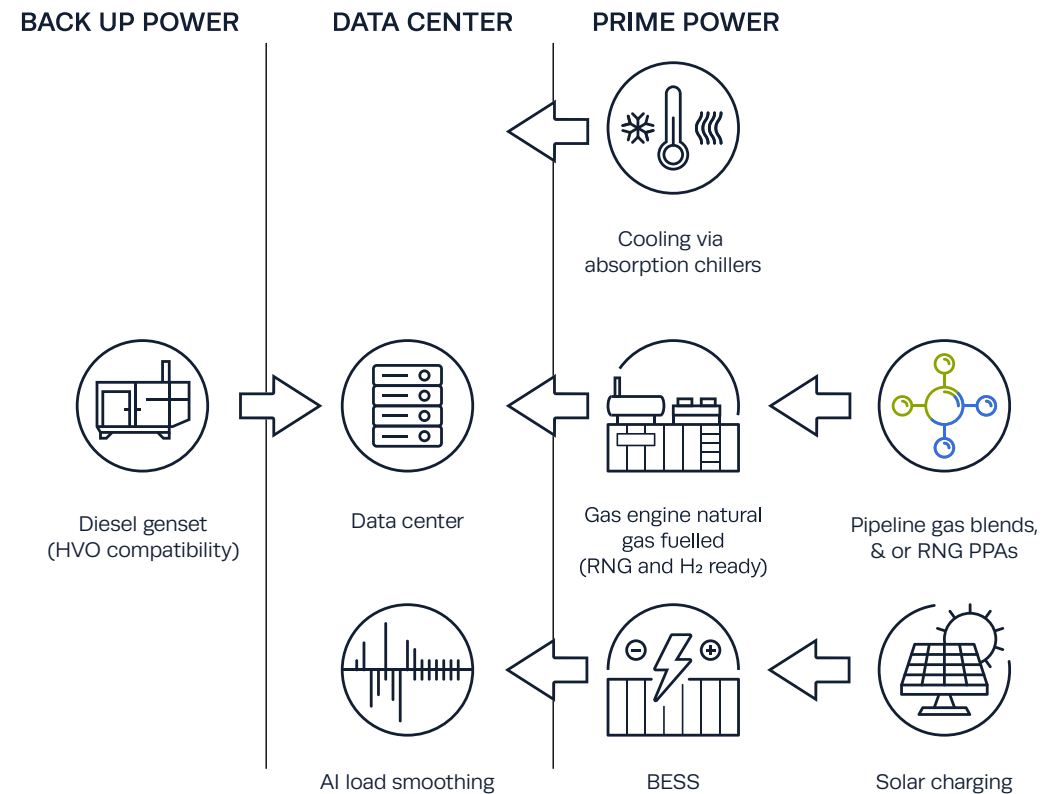


Figure 8: Integration of absorption chillers for cooling, solar power to decarbonize BESS charging and introducing blends of renewable gases such as RNG, hydrogen directly or via RGPAs.

Core Value Streams

Flexible, transition-ready power systems are positioned to access a broad range of grid and market opportunities, including:

- **Capacity markets**
Providing firm, dispatchable capacity to support system adequacy, particularly in constrained or rapidly growing regions.
- **Ancillary services**
Supplying frequency regulation, voltage control, and operating reserves through fast-responding generation and energy storage assets.
- **Demand response**
Modulating on-site generation and electrical load to relieve grid stress during peak demand periods or grid contingency events.
- **Energy arbitrage**
Optimizing imports, exports, and storage dispatch in response to real-time and forward electricity price signals.
- **Behind-the-meter optimization**
Reducing peak demand charges, improving load factor, and managing exposure to energy price volatility.

Role Within the Structured Transition Model

The importance and sophistication of value stacking increases as systems advance along the transition pathway:

Phase 1

Limited market participation, with value creation primarily driven by reliability, resiliency, and peak demand management.

Phase 2

Active participation in capacity and ancillary service markets, enabled by battery storage, hybrid generation, and advanced control platforms.

Phase 3

Fully integrated market participation, with on-site assets operating as flexible grid resources while continuing to meet core resilience and uptime requirements.

System Design Implications

To unlock these stacked value streams, systems must be designed from inception with:

- Fast-response assets, including energy storage and flexible generation
- Advanced control and dispatch platforms capable of real-time optimization
- Multi-source electrical architectures that enable seamless operation across grid, on-site generation, and storage
- Market-ready configurations, including grid-code compliance and participation eligibility for relevant market products

Three Phases: The Structured Transition Pathway

System-level benefits of biogas and biomethane integration

Hybridization is not limited to power electronics and thermal recovery. Fuel selection also plays a critical role in determining whether emissions reductions are confined to the site boundary or delivered at system scale. The World Biogas Association's Biogas: Pathways to 2030 report highlights that full global deployment of anaerobic digestion and biogas utilisation could enable a 10–11% reduction in global greenhouse gas emissions by 2030, primarily through the capture and productive use of methane from organic waste that would otherwise be released to atmosphere.

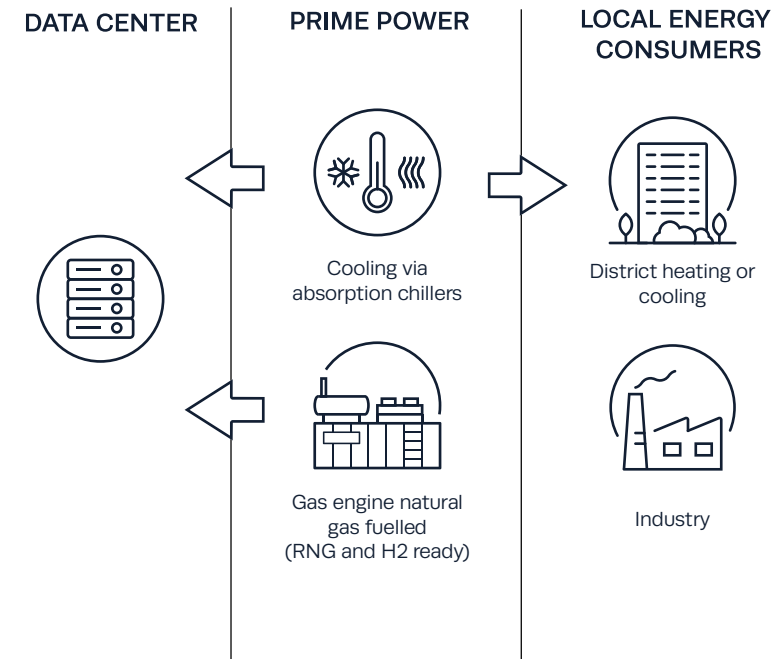
In this context, the use of biogas or RNG/biomethane-fuelled generation within data-center hybrid power systems delivers benefits analogous to gas to power and CHP. Just as CHP improves overall energy efficiency by utilizing waste heat, biogas systems improve global climate outcomes by capturing and utilising methane emissions upstream in the waste system.

The resulting carbon benefit is therefore not limited to the data center itself, but contributes to broader methane abatement and circular-economy objectives.

Data centers are uniquely positioned to support this pathway. Their large, stable, and long-term energy demand enables them to act as anchor customers for biomethane supply, supporting investment in anaerobic digestion infrastructure, renewable gas upgrading and distribution. As such, hybrid data-center developments using biogas-fuelled generation contribute not only to their own decarbonization trajectory, but to wider global emissions reductions consistent with the Pathways to 2030 framework.

© World Biogas Association (2020) Biogas: Pathways to 2030

Figure 9: Forward energy efficient thinking integrates cooling integration with absorption chillers or integrate with district energy networks, almost doubling the amount of useful energy per unit of fuel and typically has lower CO2 intensity compared to local electricity networks.



Three Phases: The Structured Transition Pathway

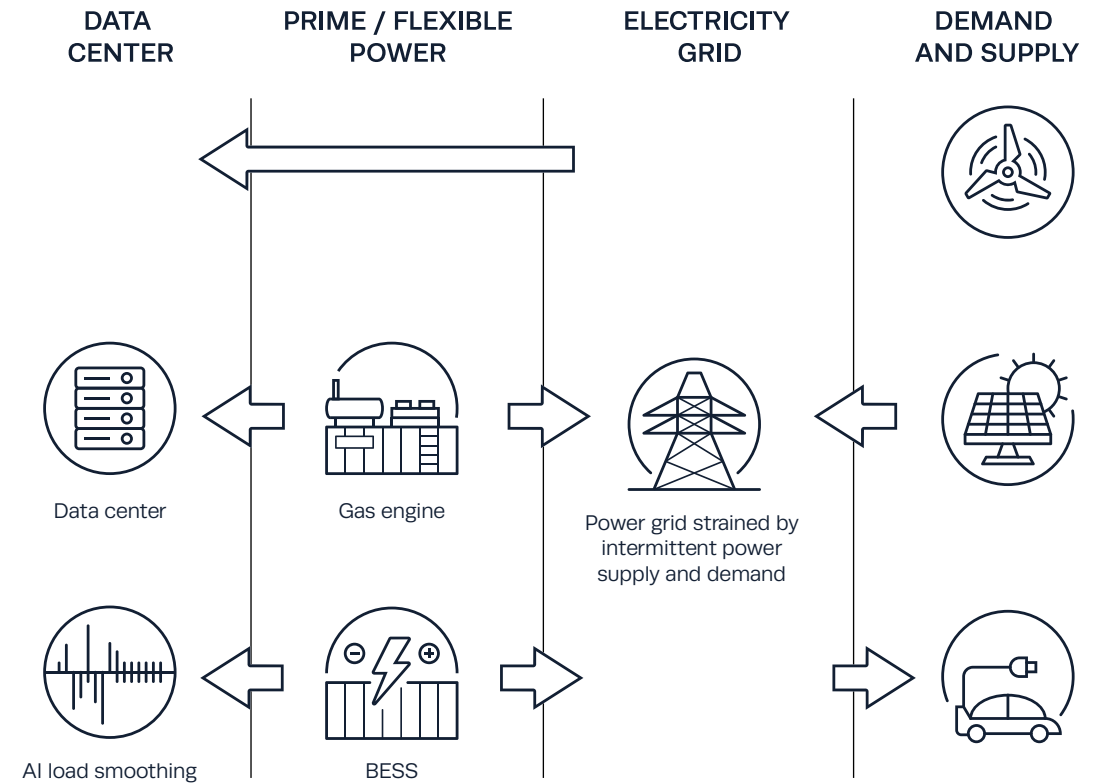
Phases 2-3:

Fuel Transition and Grid Integration (Years 7–20)

Phase 3 builds directly on the hybrid architectures established in Phase 2, shifting the decarbonization focus from system optimisation toward fuel transition and grid integration without altering core infrastructure. In this framing, AI data centers evolve from passive electricity consumers into active participants in energy-system reliability, efficiency and decarbonization.

The final phase aims to achieve deep decarbonization without stranding core infrastructure. Depending on local conditions and asset specifications, the fuel mix can shift toward renewable gases, including direct biomethane substitution. Where feasible and in line with equipment capabilities and safety codes, low carbon hydrogen blending or conversion may be pursued. The renewable supply portfolio can expand and diversify, and as interconnection constraints ease, the grid may progressively supply a greater share of low carbon energy.

Figure 10: Onsite assets can support local grids strained by intermittent demand and supply



⁷ Depending on feedstock, system configuration, and lifecycle accounting methodology, certain pathways may achieve very low or potentially net-negative lifecycle emissions; such outcomes are highly dependent on project-specific assumptions and require independent validation.

Three Phases: The Structured Transition Pathway

At this stage, on site dispatchable generation functions primarily as strategic resilience capacity, preserving uptime during grid disturbances while providing optionality in market participation. Critically, the infrastructure built in Phase 1 remains fully usable: electrical topologies accept multiple sources, engines are compatible with lower carbon fuels, and controls can orchestrate a hybrid fleet. Capital replacement cycles are aligned with technology maturity and fuel availability, allowing upgrades to be timed for both performance and cost effectiveness. Underutilized assets can be monetized through ancillary services and flexibility markets, further improving project economics. Success is defined by a lifecycle emissions trajectory aligned with a credible net zero pathway, the absence of stranded core assets and a reliability envelope that is maintained or improved. In locations with mature low carbon grids and secured renewable fuel or PPA positions, developments may launch directly at Phase 3; including fully decarbonised operational modes, while still adopting the architecture and contractual frameworks that preserve resilience and optionality.



Gas Turbines and the Structured Transition Model

Gas turbines represent a proven and widely deployed form of dispatchable generation and continue to play a central role in national power systems, industrial cogeneration and utility scale capacity additions.

They are particularly well suited to applications requiring large blocks of power, high exhaust grade heat, and sustained steady state operation. In these contexts, gas turbines offer high power density and scalability at the grid and industrial system level.

Within the Structured Transition Model, however, gas turbines are evaluated alongside other dispatchable technologies based on their operational characteristics, rather than treated as a default solution for AI data-center power.

Gas turbines are generally most efficient and economically attractive when operating close to rated output or within relatively narrow load bands. While modern turbine designs continue to improve start up times and turndown ratios, frequent cycling and extended part load operation typically impose efficiency penalties, increase thermal stress, and reduce overall utilization of installed capacity.

By contrast, the Structured Transition Model places greater emphasis on generation assets that can operate efficiently across a wide range of load conditions, accommodate frequent cycling, and migrate over time from prime power toward balancing, resilience and flexibility roles. This operational adaptability becomes increasingly important as renewable penetration rises, battery storage is integrated, and dispatchable assets shift from continuous duty toward system support functions.

This distinction does not exclude gas turbines which may remain appropriate in specific scenarios, including very large single site campuses with stable, high utilization profiles; locations where high temperature thermal output (steam) is prioritized for industrial processes; or where integration occurs primarily at the utility scale rather than within a distributed microgrid framework. In such cases, turbines can form part of the broader regional energy system within which data centers operate.

⁸IEA (2023) Electricity Grids and Secure Energy Transitions

⁹US EPA (2024) Efficient Generation: Combustion Turbine Electric Generating Units, Technical Support Document

Nuclear Power and the Structured Transition Model

Nuclear energy is frequently cited as a long-term solution for zero-carbon baseload power, including for data center campuses. While nuclear may play a significant role in national electricity systems over multi-decadal horizons, it is not treated as a core element of the Structured Transition Model.

The model is explicitly designed around near-term deployability, modular scaling, private control, and phased evolution within a typical 10–20-year data-center asset lifecycle. Current nuclear technologies, including small modular reactors, face development timelines, permitting complexity, large and indivisible capital commitments, and regulatory uncertainty that place them outside the practical decision window for most AI data center developments today. As a result, nuclear is best understood not as a transitional technology within the model, but as a potential future grid characteristic that structured, flexible power architectures are designed to accommodate should it materialise at scale.

In addition to deployment and permitting considerations, nuclear generation is inherently optimised for high-capacity-factor, steady-state operation. While modern reactor designs can

provide some degree of load-following capability, nuclear plants deliver their strongest economic and carbon performance when operated continuously rather than cycled in response to short-term load variability. This is a conclusion reflected in OECD Nuclear Energy Agency analysis, which identifies high-capacity-factor operation as the most economical and technically straightforward operating mode for nuclear power plants, with load-following introducing additional system and asset-management considerations.

AI data center demand profiles, by contrast, are increasingly dynamic, with rapid ramping, fluctuating utilization and growing interaction with variable renewable supply. This operational mismatch reinforces the need for flexible, fast-responding assets — such as engine-based generation, storage, and hybrid microgrids — to manage volatility at both site and regional levels. In this sense, the Structured Transition Model does not compete with, nor pre-empt, future nuclear deployment; it creates the flexibility envelope within which nuclear generation, where available, can operate most effectively as part of a balanced, resilient energy system.

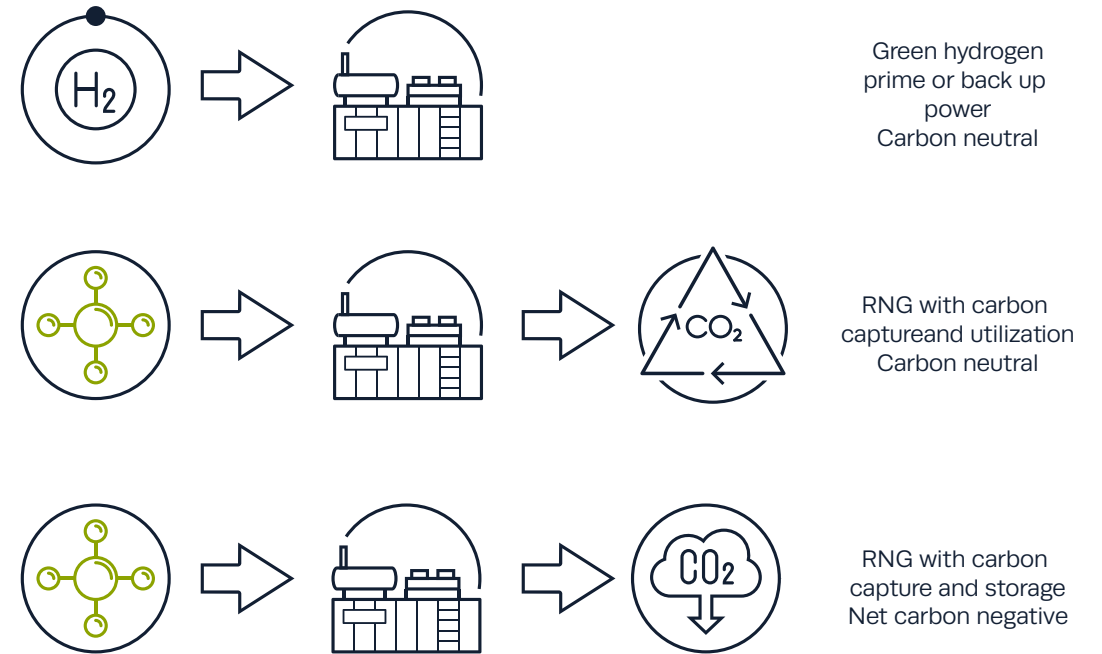
¹⁰IEA (2025) Nuclear Power and Secure Energy Transitions <https://www.iea.org/reports/nuclear-power-and-secure-energy-transitions/executive-summary>

¹¹OECD Nuclear Energy Agency (2021) Technical and Economic Aspects of Load Following with Nuclear Power Plants

4: Lifecycle Thinking vs. Day 1 Carbon

Procurement practices often optimise for first year emissions, producing designs that look favourable at commissioning but underperform over the long term. Carbon performance depends on when a project is built, how the local and regional grid evolves, the availability and cost of lower carbon fuels, the asset's dispatch profile and the timing of technology refresh cycles. A system that appears imperfect on day one can outperform alternatives over its lifecycle if it enables earlier deployment of efficient compute, supports incremental integration of renewables and storage, and avoids the need for wholesale infrastructure replacement as markets and technologies mature. The Structured Transition Model therefore evaluates total lifecycle impact rather than static design snapshots, guiding decisions toward architectures that remain adaptable, financeable and credible as decarbonisation pathways unfold. Lifecycle assessment within this model also accommodates late entry projects by back casting design choices from the chosen starting phase to ensure compatibility with prior and subsequent phases, preserving both asset value and transition credibility.

Figure 11: Deep decarbonization: onsite prime power assets can be fuelled by renewable fuels, with the integration of carbon capture and storage, facilities can deliver net carbon negativity.



¹² Depending on feedstock, system configuration, and lifecycle accounting methodology, certain pathways may achieve very low or potentially net-negative lifecycle emissions; such outcomes are highly dependent on project-specific assumptions and require independent validation.

4a: Worked Example

The modelled system includes twenty-two gas engines, with no redundancy, delivering ~98MW and operating 8,000 hours annually over a 15-year project life.

The example uses Indiana as a representative U.S. location with an average grid carbon intensity of 415gCO₂e/kWh.

Because grid carbon intensity fluctuates hourly with renewable availability and marginal dispatch decisions, using an annual average may materially underestimate the carbon advantage of CHP relative to marginal emissions, which are frequently significantly higher during periods of peak demand or low renewable output.

The following values are illustrative modeled results for comparative purposes only and should not be used for project-specific engineering, financial, or regulatory determinations.

Generation type	Fuel	CCP (0.7 COP)	CCUS	Estimated CO ₂ emissions (gCO ₂ e/kWh)	Intensity of power generation versus average local grid (%)
Electricity	Natural gas	N	N	490	+18
CHP	Natural gas	Y	N	280	-33
Electricity	Natural gas	N	Y	55	-87
CHP	Natural gas	Y	Y	40	-90
Electricity	Natural gas + 20% H ₂	N	N	455	-10
CHP	Natural gas + 20% H ₂	Y	N	260	-37
Electricity	Natural gas + 20% H ₂	N	Y	50	-88
CHP	Natural gas + 20% H ₂	Y	Y	38	-91
Electricity	Hydrogen	N	N	25	-94
Electricity	RNG	N	N	30	-93
CHP	RNG	Y	N	15	-96
Electricity	RNG	N	Y	-50	-112
CHP	RNG	Y	Y	-80	-119

Footnotes & Assumptions

A – System Boundary

All values represent well to wire lifecycle carbon intensities, including upstream fuel emissions, combustion, parasitic loads, and avoided emission crediting where applicable. Values are intended for illustrative comparison and should not be used for project finance or regulatory reporting.

B – Natural Gas Baseline

The natural gas baseline assumes reciprocating engine generation with typical U.S. upstream methane leakage (~1.4%) and a heat rate consistent with commercial prime power engines. This yields ~450 gCO_{2e}/kWh, consistent with published LCA ranges for engine-based power generation.

C – CHP Calculations

CHP intensities use a 0.7 COP absorption chiller, allocating carbon reductions through displaced grid electricity for cooling. Single-effect absorption chillers can operate between 0.6 and 0.8. The CHP benefit varies with thermal load; a moderate, defensible value of 320 gCO_{2e}/kWh is used.

D – CCUS Efficiency and Parasitics

CCUS scenarios assume 90–95% capture of stack CO₂ and a parasitic electrical load of 6–8%, including capture, compression, and ancillary equipment. Residual emissions arise from incomplete capture and upstream methane emissions from fuel supply.

E – Hydrogen Blends

The 20% hydrogen blend assumes 20% energy substitution, which corresponds to approximately 18–20% reduction in lifecycle carbon intensity. Hydrogen production emissions are allocated to natural gas share; H₂ is assumed to be low-carbon.

F – Green Hydrogen Assumption

100% hydrogen scenarios assume renewable (green) hydrogen produced from electrolysis powered by renewable electricity. Emissions are attributed to electrolyser manufacturing and renewable infrastructure, resulting in a residual ~10 gCO_{2e}/kWh footprint.

G – RNG and Negative Emissions

RNG scenarios assume waste-derived biomethane, with biogenic carbon from the short-term carbon cycle where avoided methane emissions (e.g., manure, wastewater, landfill gas) provide significant negative carbon credits under standard lifecycle analysis methods. Values of –50 to –100 gCO_{2e}/kWh reflect typical avoided methane crediting without reaching extreme best-case assumptions.

H – RNG + CCUS

CCUS applied to waste-derived RNG captures biogenic CO₂, adding to the negative avoided methane credit. Conservative combined values of –150 to –200 gCO_{2e}/kWh are used, although literature indicates some pathways can achieve larger negative intensities depending on feedstock and allocation methodology.

I – Impact of CHP on marginal emissions

This comparison uses average grid carbon intensity. Because grid emissions fluctuate hourly, annual averages may significantly understate the carbon advantage of CHP relative to marginal grid emissions, particularly during periods of low renewable availability or peak demand.

J – Illustrative Values Only

All values are indicative, based on representative technology and LCA parameters. Actual performance depends on equipment selection, site conditions, fuel supply, waste stream characteristics, thermal integration, and CCUS system design.

Assumes the same CCUS capture rate and parasitic load as fossil CCUS cases.

4b: Translating the Structured Transition Model into Data Center Performance Metrics

While the Structured Transition Model is presented as an infrastructure and lifecycle design framework, its relevance ultimately rests on how it influences the core performance metrics used by data center operators. Decisions are rarely made on energy systems in isolation; they are evaluated through their impact on efficiency, reliability, and carbon intensity at the facility level.

In conventional approaches, these metrics are largely determined by external factors—most notably the carbon intensity of the grid and the efficiency of electrically driven cooling systems. The Structured Transition Model introduces a different dynamic. By integrating dispatchable generation, thermal recovery, storage, and flexible fuels within a single architecture, it allows operators to actively shape these outcomes over time rather than passively inherit them.

Several established data center metrics provide a useful lens through which to evaluate this impact.

Power Usage Effectiveness (PUE) remains the most widely used measure of facility efficiency, expressing the ratio of total site energy to IT energy. In grid-only configurations, PUE is heavily influenced by electrically driven cooling demand and grid losses. Within the Structured Transition Model, particularly in Phase 2, the integration of combined cooling, heat and power (CCHP) allows thermal energy from prime movers to be

repurposed for cooling via absorption chillers. This reduces reliance on electrically driven cooling and improves overall system efficiency, lowering effective PUE.

Closely related is Data Center Infrastructure Efficiency (DCiE), the inverse of PUE, which provides a direct measure of how much input energy is delivered to IT load. Improvements in PUE under hybrid and thermally integrated configurations translate directly into higher DCiE, reflecting a greater proportion of useful energy.

Cooling Load Factor (CLF), defined as the ratio of cooling demand to IT load, is a less commonly cited but increasingly important metric as AI workloads drive higher thermal densities. In traditional architectures, cooling demand is met almost entirely through electrically powered systems, resulting in relatively high CLF values. By contrast, the Structured Transition Model enables a portion of this cooling demand to be met through recovered heat, effectively reducing the electrical burden associated with cooling and improving overall system balance.

From a carbon perspective, Carbon Usage Effectiveness (CUE) provides a direct measure of emissions per unit of IT energy. Grid-only designs can perform well where low-carbon electricity is available, but remain exposed to marginal grid intensity and temporal variability. Diesel-based backup systems, when dispatched,

carry a high carbon intensity, an impact which can be mitigated with transition to HVO fuel. The Structured Transition Model takes a different approach, delivering a declining carbon trajectory over time through efficiency gains, hybridization, and fuel transition. As renewable electricity, renewable gases, and storage are integrated, CUE improves progressively rather than being fixed at commissioning.

Energy Reuse Effectiveness (ERE) captures the benefit of recovering and reusing energy that would otherwise be rejected as waste heat. This is an area where the Structured Transition Model offers a clear structural advantage. Through CHP and CCHP integration, heat can be reused for cooling or exported to adjacent processes or district energy systems, materially improving overall energy utilisation and reducing net energy demand.

Taken together, these metrics illustrate a broader shift in how performance is defined. Traditional models tend to optimise for a static point—typically first-year efficiency or emissions. The Structured Transition Model instead enables continuous optimisation, with performance improving as additional layers of the system are integrated and as external conditions evolve.

The table below summarises how these metrics typically compare across common power strategies.

4b: Translating the Structured Transition Model into Data Center Performance Metrics

In practice, the significance of these improvements lies not only in headline efficiency gains, but in control. By internalizing key elements of power and thermal management, operators gain the ability to optimise across reliability, cost, and carbon simultaneously, rather than trading one against another. This becomes increasingly valuable as AI-driven loads introduce greater variability and as expectations around both uptime and sustainability continue to tighten.

In this sense, the Structured Transition Model does not redefine data center performance metrics—it changes how they are achieved and how they evolve over time.

Metric	Grid-Only (PPA / Utility Supply)	Diesel Backup Model	Structured Transition Model
Power Usage Effectiveness (PUE)	Dependent on grid and electric cooling efficiency	Typically, higher when backup systems are used	Improved through CHP/CCHP and reduced electrical cooling demand
Data Center Infrastructure Efficiency (DCiE)	Moderate	Lower when backup generation is active	Higher due to improved overall system efficiency
Cooling Load Factor (CLF)	High (electrically driven cooling)	High	Reduced through thermal recovery and absorption cooling
Carbon Usage Effectiveness	Dependent on grid carbon intensity	High when dispatched	Declining over lifecycle through hybridisation and fuel transition
Reliability (Availability)	Grid dependent	High for backup, limited for prime operation	Very high through modular, redundant generation architecture

Avoiding Stranded Assets Through Transition Ready Design

For developers and investors, one of the most material risks in modern data center infrastructure is the potential for stranded assets: equipment or systems that become economically or operationally obsolete long before the end of their technical life. This risk is heightened in AI driven facilities, where rapid growth in demand coincides with evolving fuel markets, tightening carbon policy and uncertain grid conditions. Traditional design approaches that optimise for a single end state or a narrow emissions snapshot can lock in technologies that later require costly replacement or cannot adapt to new operating realities.

The Structured Transition Model is explicitly designed to mitigate this risk. By prioritizing fuel flexible generation

assets, multi source electrical topologies and modular plant design, the model ensures that early stage infrastructure remains usable throughout the asset lifecycle. Technologies deployed in Phase 1 continue to provide resilience in later phases, even as their operational role shifts from prime supply to balancing and emergency capacity. Hybridisation introduced in Phase 2 extends asset relevance rather than displacing it, while Phase 3 synchronizes capital replacement with market maturity and fuel availability rather than forcing premature upgrades.

From an investment perspective, this approach protects return profiles by reducing both regulatory and technical obsolescence risk. It limits the exposure

to sudden shifts in carbon pricing, fuel supply volatility or policy constraints, and enables projects to evolve with the decarbonisation trajectory of the local grid. Rather than investing in rigid, single purpose assets that may be misaligned with future requirements, the model encourages architectures that remain economically useful, operationally resilient and compliant with emerging standards over their full design life. In this way, transition ready design becomes a form of capital protection as well as an enabler of long term decarbonisation.

6:

Political Risk & Infrastructure Longevity

Energy infrastructure deployed to support AI data centers typically operates for 15 years or more, often extending beyond this horizon as engines and other equipment are overhauled, upgraded, or repowered. Over such a timeframe, the surrounding political and regulatory environment is almost certain to shift. Government priorities, environmental policy, electricity market structures and permitting regimes can all evolve materially, reshaping how assets are dispatched, valued and integrated with the grid. Additionally, international and geopolitical events can affect gas and electricity prices, switching the economic priorities of the host country. Unlike short lived compute hardware, power systems must therefore be designed to operate through multiple cycles of political and regulatory change.

Recent U.S. Department of Energy analysis underscores that without new firm capacity and changes to traditional planning approaches, most regions will face unacceptable reliability risks within five years—driven largely by AI-related load growth from data centers.

The Policy Cycle

Across a 15 Year Asset Life

In markets such as the United States, energy related policy evolves at several overlapping levels; federal executive direction, congressional legislation, agency rulemaking, state energy policy, and regional electricity market design. Each layer operates on its own cadence. Over a typical 15 year asset life, an AI data center power system may experience:

- Three to four shifts in federal executive priorities
- Two or more rounds of major federal fiscal or climate related legislation
- Several rulemaking cycles affecting emissions standards, market access or environmental compliance
- Multiple adjustments to state level renewable or low carbon policies
- One or two significant reforms to electricity market design or interconnection frameworks

Taken together, a reasonable expectation is that 8–15 meaningful policy interventions can occur during the life of a local distributed energy system. These interventions can influence emissions limits, carbon pricing mechanisms, fuel incentive structures, permitting requirements, grid service participation, and the economics of hybrid and resilient microgrids. Importantly, not all policy changes generate downside risk: flexible distributed generation, thermal integration and hybrid system designs often unlock new revenue streams or enhanced operating roles under evolving market rules.

¹³ U.S. Department of Energy (2025) – Evaluating U.S. Grid Reliability and Security

Designing for Policy Uncertainty

The Structured Transition Model does not attempt to predict the policy landscape of the 2030s or 2040s. Instead, it acknowledges uncertainty as an inherent feature of long lived infrastructure. By assuming that policy change is inevitable, the model focuses on the technical and architectural characteristics that preserve operational and financial viability across multiple regulatory regimes.

This approach prioritises assets and configurations that can adapt, rather than those optimised only for current incentives or constraints.

Key characteristics include:

- **Deployability today:** Technologies that can meet near term reliability needs without waiting for unproven breakthroughs.
- **High electrical efficiency:** Providing a natural hedge against future carbon pricing or tightening emissions standards.
- **Combined heat and power integration:** Capturing heat for cooling or external use, improving lifecycle emissions performance and future proofing against stricter environmental frameworks.
- **Fuel flexibility:** Supporting pathways from natural gas to biomethane, hydrogen blends and ultimately zero carbon fuels as policy and market conditions evolve.
- **Hybrid system architecture:** Enabling the integration of battery storage and renewable generation as electricity market rules shift toward flexibility and decarbonization.
- **Modularity:** Allowing targeted upgrades, repowering or fuel transitions without wholesale replacement of core infrastructure.

6: Political Risk & Infrastructure Longevity

Political Risk as a Design Parameter

Within this framework, political change is treated not solely as a risk factor but as a design parameter for long lived power systems. Infrastructure that can evolve alongside policy remains valuable, resilient and compliant even as regulations change. Conversely, systems locked into narrow fuel pathways, rigid operating modes or incompatible architectures face a higher likelihood of premature obsolescence.

For digital infrastructure developers, investors and operators, the central insight is clear:

The greatest long term exposure is not regulatory change itself, it is deploying infrastructure that cannot adapt to it.

The Structured Transition Model therefore positions policy uncertainty as a catalyst for flexible, transition ready design. By delivering reliability today while preserving pathways for future decarbonisation and regulatory alignment, the model protects asset value and supports continued operational viability across the full lifecycle of AI driven data center development.



Implications for Developers, Investors & Policymakers

For developers, the model argues for specifying generation assets that are certified or upgradeable for renewable gases or hydrogen blending as applicable, and for designing electrical architectures that accept multiple sources from the outset. Sites should be conceived as future hybrid microgrids even if initial operation is simpler, so that integrating storage, renewables, heat recovery or grid services later is straightforward. Practically, this means selecting a starting phase that reflects local constraints and opportunities - including beginning at a hybrid or fully decarbonized state - while ensuring forward and backward compatibility across phases.

For investors, the emphasis shifts from initial carbon scores to transition readiness. Projects that exhibit retrofit agility—through modular plant design, fuel flexible equipment and advanced control layers—reduce deployment risk and enhance long term value. Optionality

itself becomes an asset characteristic, hedging regulatory uncertainty, market volatility and technology risk.

For policymakers, the approach supports regulation that values credible, staged transition pathways over rigid technology mandates. Interconnection frameworks and market rules that reward flexible, hybridised infrastructure will accelerate deployment while preserving reliability. Recognizing dispatchable low carbon fuels as practical enablers helps bridge the near term gap between today's reliability needs and tomorrow's decarbonized grid.

Conclusion:

From Static Design to Dynamic Infrastructure

AI data centers are not static buildings; they are long lived energy systems operating within an evolving grid and fuel landscape. The central choice is not between immediate perfection and short term expediency, but between ad hoc buildouts and structured evolution. By prioritizing reliability today, enabling hybridisation and optimization as markets develop and technologies mature, and allowing for deep decarbonization through planned fuel transition and grid integration, developers can meet surging compute demand without accruing technical debt or constraining future options. Crucially, there is no single mandated starting line: some projects will begin at the reliability foundation, others at advanced hybridization, and an increasing number, where market and grid conditions permit, will proceed directly to fully decarbonized architectures while preserving the same principles of resilience and optionality. For AI-driven growth, sustainability is no longer defined by what is installed on day one, but by how fast reliable capacity arrives, and how effectively that infrastructure evolves.

¹⁴ Depending on feedstock, system configuration, and lifecycle accounting methodology, certain pathways may achieve very low or potentially net-negative lifecycle emissions; such outcomes are highly dependent on project-specific assumptions and require independent validation.



About the Author

Alex Marshall serves as Group Business Development and Marketing Director for Clarke Energy a Rehlko Company, leading strategic growth across distributed power, hybrid energy systems and clean fuel integration. Beyond his corporate leadership, he is Vice President of the COGEN World Coalition, plays an active role in the iMasons Climate Accord Energy Working Group, and sits on the Council of the World Biogas Association, contributing to industry standards and global decarbonisation initiatives. Alex is widely regarded for his ability to connect technical insight with commercial and policy realities, helping developers, investors and governments build energy systems that balance reliability, cost and carbon while remaining flexible enough to evolve.

Contributions & Credits

John Dangles – Sr Content Manager - Rehlko

Ryan DeLoge – ESG Manager – Rehlko

Louis Liu – Sr Project Engineer – Sustainability – Rehlko

Dalia Si Ahmed – Group Applications Manager – Hydrogen – Clarke Energy

Justin Lenoff – Group Applications Manager – Microgrids – Clarke Energy

Glossary

ASTM D975

The ASTM standard specification that defines the required properties and allowable limits for diesel fuel oils used in compression-ignition engines.

Reciprocating Gas Engines

Modular, fast responding on site gas-fuelled generators that provide high reliability and strong part load performance. Central to Phase 1 and the Structured Transition Model.

CHP / CCHP

Combined Heat and Power (CHP) and Combined Cooling, Heat and Power (CCHP) reuse engine heat for heating or cooling, improving efficiency and reducing carbon intensity.

BESS (Battery Energy Storage System)

Short duration storage used to smooth AI load variability, reduce peaks, and participate in grid service markets.

HVO (Hydrotreated Vegetable Oil)

A renewable diesel fuel made by hydrotreating vegetable oils or waste fats, producing a drop in diesel with very low sulfur and good cold weather performance.

Microgrid Controls

Software and control platforms that coordinate engines, renewables, storage, and the grid to optimize cost, reliability, and carbon. Essential for value stacking.

RNG (Renewable Natural Gas)

Waste derived biomethane with potentially negative lifecycle emissions due to avoided methane release. Used in Phase 3 fuel transition. Also called biomethane.

Hydrogen (H₂)

Low carbon fuel used as a blend or pure fuel in later stages. In the model, hydrogen refers to renewable ("green") hydrogen.

CCUS (Carbon Capture, Utilization & Storage)

Captures up to ~90–95% of engine CO₂ emissions. When combined with RNG, can produce strongly negative lifecycle carbon performance.

PUE (Power Usage Effectiveness)

Efficiency ratio of total facility energy to IT energy. Lower values are better; CCHP reduces PUE by shifting cooling to thermal sources.

DCiE (Data Center Infrastructure Efficiency)

Inverse of PUE. Higher values indicate a greater share of input energy is delivered to IT equipment. Improved under integrated hybrid systems.

CUE (Carbon Usage Effectiveness)

Carbon emissions per unit of IT energy. Decreases across Phases 2 and 3 as the system hybridizes, improves efficiency, and transitions fuels.

Value Stacking

Revenue streams from grid services such as capacity markets, frequency regulation, demand response, and energy arbitrage. Strongest in Phases 2 and 3.

Stranded Asset Risk

Risk of infrastructure becoming unusable due to carbon policy, grid constraints, or technology shifts. The model mitigates this via modularity and fuel flexibility.

Transition Ready Design

Designing power architecture so today's assets remain compatible with future fuels, hybridization, grid participation, and decarbonization targets.

References

Lenoff, J. (2025) Powering AI Data Centers: Hybrid Energy Strategies for Extreme Volatility, Clarke Energy, Rehlko.

Cogen World Coalition (2025) CHP – A Key Enabling Technology for Ensuring Data Center Development is Efficient, Cost-Effective and Resilient, https://cogenworld.org/wp-content/uploads/2025/12/Data_Center_Cogeneration_Opportunities_CWC_2025_12_04_1_.pdf

International Energy Agency (2023) Electricity Grids and Secure Energy Transitions <https://www.iea.org/reports/electricity-grids-and-secure-energy-transitions/executive-summary>

International Energy Agency (2025) Energy and AI <https://www.iea.org/reports/energy-and-ai>

International Energy Agency (2025) Nuclear Power and Secure Energy Transitions
IEA (2025) Nuclear Power and Secure Energy Transitions <https://www.iea.org/reports/nuclear-power-and-secure-energy-transitions/executive-summary>

iMasons (2025) Climate Accord Maturity Model, A Guide to Emissions Reduction in Digital Infrastructure; Practical Decarbonization in Power, Equipment & Materials; https://climateaccord.org/wp-content/uploads/2025/02/iCA_Maturity_Model_Report_2025_v1.1.pdf

OECD Nuclear Energy Agency (2021) Technical and Economic Aspects of Load Following with Nuclear Power Plants https://www.oecd.org/en/publications/technical-and-economic-aspects-of-load-following-with-nuclear-power-plants_29e7df00-en.html

U.S. Department of Energy (2025) – Evaluating U.S. Grid Reliability and Security <https://www.energy.gov/sites/default/files/2025-07/DOE%20Final%20EO%20Report%20%28FINAL%20JULY%207%29.pdf>

U.S. EPA (2024) Efficient Generation: Combustion Turbine Electric Generating Units, Technical Support Document. <https://www.epa.gov/system/files/documents/2024-04/tsd-efficient-generation-combustion-turbine-egus-april-2024.pdf>

World Biogas Association (2020) Biogas: Pathways to 2030 <https://www.worldbiogasassociation.org/biogas-pathways-to-2030-report>

World Economic Forum (2025-6) – Data Centers as Strategic Energy Infrastructure

Third-party data, studies, and reports are referenced for illustrative purposes only. Rehlko has not independently verified all third-party assumptions or methodologies.