WÄRTSILÄ

Unlock power. Unleash growth.

How data centres can create tomorrow's power system flexibility

A study by Wärtsilä Energy



Contents

1. Powering the age of AI 5
1.1. Where digital and energy infrastructure converge 5
1.2. The growing power demand of the digital economy 5

2. Solving the energy grid bottleneck 6
2.1. How energy grid constraints are impacting growth 6
2.2. The rise of onsite power generation 6

3. Powering the off-grid data centre buildout 7
3.1. Making the right power choices for data centres 7
3.2. The scale of power capacity required 8
3.3. The economics of technology choice 9
3.4. Minimising data centre land use 10
3.5. Minimising data centre water use 11

4. Enabling the path towards a 100% renewable energy future 12
4.1. Transitioning beyond off-grid power 12
4.2. Unlocking a new source of grid flexibility 15
4.3. Reimagining data centre backup power 17
4.4. Turning waste heat into value 17

5. From servers to system resources 18
5.1. The limits of flexible compute 18
5.2. Harnessing waste heat from data centres 18

6. Pathways for net zero data centres 19

7. Conclusion 21

Some images in this report have been created or modified using artificial intelligence (AI).

WÄRTSILÄ



> Developers now have an opportunity to install technology that supports the wider energy system for decades to come.

Powering the digital age, enabling the energy transition

The rapid growth of AI and data centres has raised concerns about electricity demand, constrained grids and the pace of decarbonisation. These pressures are real, but viewing data centres as a growing burden on power systems risks overlooking the opportunity they represent.

As developers consider the best ways to power data centres – and turn to off-grid energy solutions to overcome grid delays – they have an opportunity to install technology that supports the wider energy system for decades to come.

Investing in flexible technologies, such as medium-speed engines, to power data centres today could help lay the foundation for the transition towards a 100% renewable energy future. As data centres connect to the grid, these technologies could transition from primary power sources into flexible balancing capacity.

We need to think long-term. Flexibility is essential to realising the lowest cost, lowest emissions and least land-intensive transition to more sustainable power.

This report explores what that opportunity could look like at scale. It examines the technologies, infrastructure, investment and resources required if data centres increasingly generate their own power, as well as the potential value these assets could provide to the wider electricity system over time.

Our study considers a scenario in which 55 GW of onsite-powered data centre capacity is deployed globally by 2035.

Once grid-connected, mobilising this infrastructure for system flexibility could avoid up to EUR 200 billion in additional balancing-capacity costs over 20 years. That potential demands a broader conversation about how private investment in digital infrastructure can support shared energy goals and what needs to happen now to make this possible.

Developers need to invest in technology that can adapt as energy systems change. Policymakers, regulators and grid operators should recognise and reward the value that flexible assets can provide, playing different roles for key digital infrastructure and energy systems.

With foresight and coordinated action, the investment powering the digital age presents an opportunity to help accelerate the energy transition.

Anders Lindberg,
President, Wärtsilä Energy and
Executive Vice President, Wärtsilä

Summary

Artificial intelligence, digitalisation and the energy transition have become increasingly interconnected. Data centres provide the critical infrastructure underpinning the digital economy, while AI and digital technologies are themselves important enablers of more efficient, flexible and renewable power systems.

Traditionally, data centres have relied on electricity supplied through the grid, with power generated elsewhere in the energy system (offsite). Today, grid constraints are delaying the connection of new digital infrastructure and the renewable energy capacity required to power them. In response, data centre developers are investing directly in ways to power data centres off-grid, onsite.

Onsite power – electricity generated at or near the data centre rather than supplied entirely through the grid – is expected to play a significant role in new data centre development over the coming decade. This study examines a scenario in which approximately 55 GW of data centre capacity is powered onsite – and off-grid – by 2035. For scale, 55 GW represents power capacity equivalent to the peak electricity demand of the United Kingdom. At this scale, technology selection becomes significant not only for individual projects, but also for the wider energy system. Differences in system sizing, efficiency and operating costs could translate into **EUR 196 billion** in cumulative cost differences over 20 years between the solutions assessed in this report.

Technology choices also shape the land and water footprint of data centres, helping to mitigate the impact of development on people and communities. More modular, reliable and flexible power solutions – such as Wärtsilä's medium-speed engines – can reduce the need for large-scale power generation and energy storage, requiring less infrastructure and land. Of the technologies assessed in this study, the one with the lowest water consumption consumes 99.97% less water than the highest alternative, highlighting the importance of technology selection in minimising the amount of water that data centres use.

The value of this infrastructure can extend far beyond powering data centres independently from the grid. When connected to the wider energy system, onsite power generation and energy storage can also help local and regional grids by keeping energy supply stable and reliable.

Mobilising 55 GW of onsite power could avoid up to **EUR 200 billion** in additional balancing-capacity costs over 20 years, while using waste heat recovered from data centres could generate more than **EUR 180 billion** in net economic value in Europe over the same period.

Together, these opportunities could create close to EUR 576 billion in value over 20 years. This is in addition to wider benefits, including reducing the amount of land and water needed, avoided excess grid investment and enabling greater integration of renewable energy.

These solutions provide a blueprint for solving the data centre power challenge. The next questions to address are how to prioritise grid connections for the right assets and how to work with data centre developers to maximise the value their infrastructure can provide to the wider energy system.

With the right technology choices, market incentives and grid arrangements, data centres can evolve from large electricity consumers into strategic energy system assets, supporting the transition towards a 100% renewable energy future.

Types of technologies

Combined-cycle gas turbine (CCGT):
Gas turbine system that recovers exhaust heat to generate additional power.

Open-cycle gas turbine (OCGT):
Gas turbine that generates power in a single cycle without further use of the exhaust heat.

Aeroderivative gas turbine (AERO):
Aircraft-derived gas turbine adapted for stationary power generation.

Medium-speed reciprocating internal combustion engine (MSE):
Large reciprocating engine operating at medium rotational speeds.

High-speed reciprocating internal combustion engine (HSE):
Smaller reciprocating engine operating at high rotational speeds.

Chapter 1

Powering the age of AI



1.1. Where digital and energy infrastructure converge

The world has entered the age of artificial intelligence (AI). AI, cloud computing and digital services are rapidly becoming embedded in everyday life, enabling everything from online services and communications to advanced manufacturing, automation, scientific research, healthcare and business operations. Data centres provide the critical infrastructure underpinning this digital transformation and are increasingly being recognised as vital national infrastructure, supporting economic competitiveness, productivity growth, technological leadership, national security and the delivery of essential public and private services.

At the same time, AI and digitalisation are critical enablers of the energy transition. They can improve energy efficiency and forecasting, enable smarter grid management and system planning, optimise electricity consumption and generation, strengthen system flexibility, and support the integration of a growing share of variable renewable generation. [1].

The rapid growth of AI, digitalisation and the transformation of the energy system are increasingly interconnected and mutually reinforcing. Digital infrastructure depends on access to reliable and affordable electricity, while increasingly complex energy systems depend on AI and digital capabilities to operate efficiently, flexibly and reliably. Together, these increasingly interconnected developments are shaping a new era of digital infrastructure.

1.2. The growing power demand of the digital economy

Global electricity demand is projected to increase by approximately 35% by 2035, driven by the electrification of transport, industry and heating, alongside the rapid expansion of digital infrastructure [2,3]. The growth of AI, cloud computing and digital services is creating unprecedented demand for computing infrastructure and the electricity required to power it. As a result, data centres are emerging as one of the most significant sources of new power demand, accounting for close to 10% of global electricity demand growth between 2024 and 2030 [2].

According to the International Energy Agency (IEA), global installed data centre capacity could increase from approximately 100 GW in 2024 to nearly 280 GW by 2035 [1]. While the majority of this capacity is concentrated in the United States, China and Europe, growth is increasingly spreading to new regions. Expanding AI deployment, data sovereignty requirements and demand for low-latency digital services are driving data centre investment across the globe [4,5].

> The rapid growth of AI, digitalisation and the transformation of the energy system are increasingly interconnected and mutually reinforcing.

Chapter 2

Solving the energy grid bottleneck



2.1. How energy grid constraints are impacting growth

Electricity grid constraints are limiting the growth of both electricity demand and supply. Across many regions, transmission networks, substations and transformers are struggling to accommodate rapidly rising loads. Grid expansion is also being held back by long lead times for key equipment. These problems are unlikely to ease before 2030 [1,6]. As a result, data centres face long and uncertain connection timelines, while growing demand from industrial expansion and electrification continues to place additional pressure on networks that are already constrained.

Grid limitations are also slowing the connection of new electricity generation capacity. Wind and solar projects face lengthy interconnection queues and network bottlenecks. Additionally, underinvestment in flexible balancing capacity is making it more difficult for power systems to make the most of new variable renewable generation. Together, these constraints are delaying the expansion of the generation capacity needed to meet rising electricity demand [3,7].

These energy grid constraints are having a profound impact on both digital growth and the transition towards a 100% renewable energy future. Access to timely and reliable grid capacity is therefore determining where and how quickly data centres, industrial facilities and renewable energy projects can be developed.

2.2. The rise of onsite power generation

Data centre developers are increasingly looking for ways to bypass growing grid connection queues. Generating electricity at or near the data centre, onsite, can provide projects with power sooner, helping to accelerate deployment. Rather than waiting several years for sufficient grid capacity, dedicated onsite primary power solutions allow facilities to be commissioned, off-grid, independently of constrained electricity networks. This makes data centres unique among major new sources of electricity demand, as some developers are investing directly in their own power generation rather than relying on a grid connection. By operating off-grid, data centres can reduce the time needed to secure power, avoid adding demand to already constrained networks and limit the impact of rapid data centre growth on electricity prices.

As of March 2026, approximately 2 GW of operational data centres rely on onsite power generation, with a further 7.5 GW under construction [8]. Beyond this near-term pipeline, around 104.5 GW of announced projects incorporating onsite generation – using gas as a fuel – are at various stages of development [8]. While this development pipeline is substantial, uncertainty remains over how many projects will ultimately materialise. Drawing on third-party forecasts and assessing the maturity of announced projects, including their development and permitting status, we define a scenario in which up to 55 GW of onsite-powered data centre capacity could be commissioned by 2035 [6,8,9,10,11,13].



Chapter 3

Powering the off-grid data centre buildout

3.1.

Making the right power choices for data centres

Given the potential scale of onsite primary power solutions to support off-grid data centres, it is important to make sure that technology choices are informed by their impact on natural resources, as well as how these assets could ultimately support the wider energy transition.

Wärtsilä's recent study, *Power solutions for off-grid data centres,* assessed alternative onsite primary power solutions at the individual project level, showing how technology choices can materially affect installed generating capacity, operating costs and natural resource requirements. Based on a scenario in which approximately 55 GW of off-grid data centre capacity is powered onsite by 2035, this study assesses the scale of infrastructure investment and resources required to support this growth.

The study also considers the longer-term value of this infrastructure. As data centres that initially rely on onsite power progressively secure grid connections, both the data centres and their onsite power systems could become important sources of flexibility and offer wider value for the energy system. The study therefore examines how these assets could support grid balancing, capacity adequacy and renewable integration, relieve pressure on constrained grids, reduce electricity costs and system emissions, and improve energy efficiency while minimising natural resource requirements.

3.2.

The scale of power capacity required

As shown in Figure 1, technology choice has a significant impact on the total generation and energy storage capacity required to reliably supply 55 GW of global onsite-powered data centre capacity by 2035. Technologies that offer greater modularity, flexibility and availability require less installed capacity per megawatt of data centre capacity served, reducing the need for oversizing – the additional capacity installed above the data centre's rated capacity to ensure a reliable supply.

The analysis shows that widespread deployment of **engine-based solutions would require substantially less installed flexible power generation and energy storage capacity than turbine-based alternatives to serve the same level of global onsite-powered data centre demand**. This reduces the overall scale of power infrastructure required while maintaining the reliability standards needed by data centres.

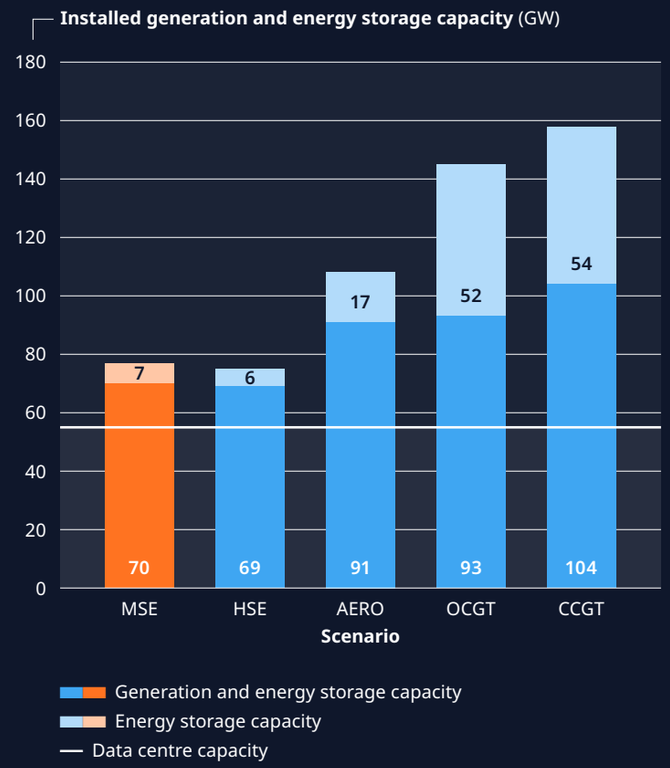
These differences are particularly important given growing supply constraints across power generation, electrical equipment and broader infrastructure supply chains. As demand for onsite power accelerates, technologies that most efficiently power data centres can reduce equipment needs, raw material requirements and exposure to manufacturing and supply-chain constraints. Technology selection therefore affects not only project economics, but the scale and speed of global onsite-powered data centre deployment.

The degree of oversizing also has important implications for the total cost and physical footprint of global onsite-powered data centre deployment. Higher installed power generation and energy storage capacity increases capital investment, fixed operating and maintenance costs, and land requirements for the same amount of data centre capacity served. These implications for costs and land use are examined in the following sections.

> The same 55 GW of data centre demand could require around 75 GW to almost 160 GW of onsite generation and energy storage capacity, depending on the technology choice.

Figure 1

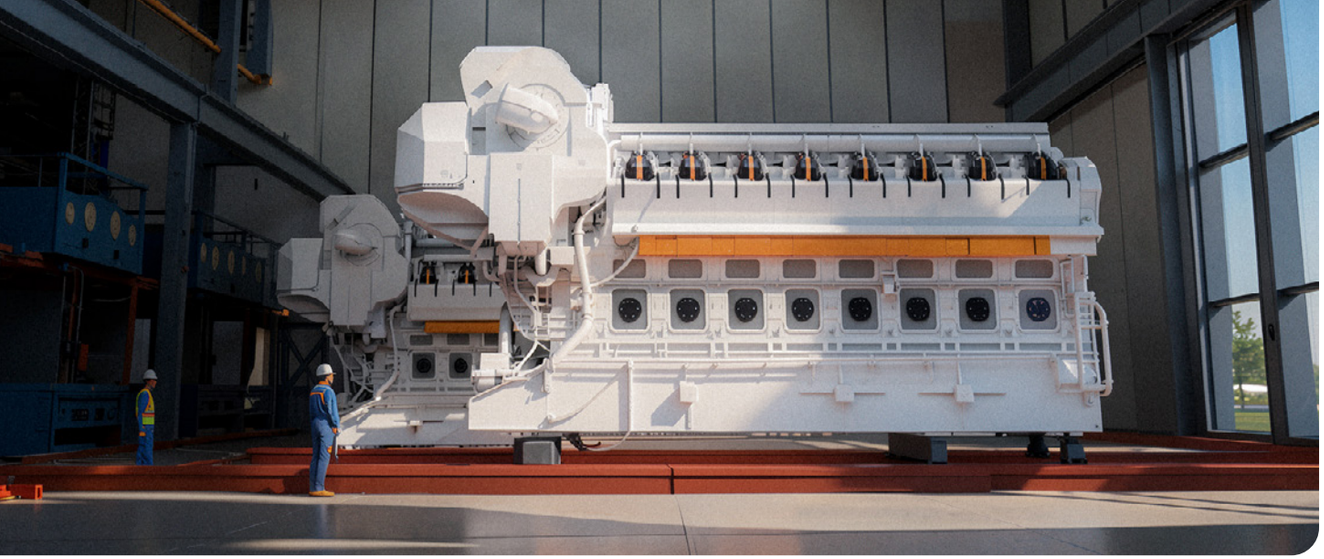
Engine-based solutions require up to **50% less** installed capacity to power 55 GW of data centres



Smaller capacity requirements reduce infrastructure needs for the same level of power reliability.

Note: Project-level sizing from the Power Solutions for Off-grid data centres study is scaled to an assumed global onsite-powered data centre capacity of 55 GW. Single-technology scenarios are illustrative, as actual deployment will involve a mix of technologies. The scenarios provide indicative lower and upper bounds for infrastructure requirements Source: Wärtsilä, 2026.

WÄRTSILÄ



3.3.

The economics of technology choice

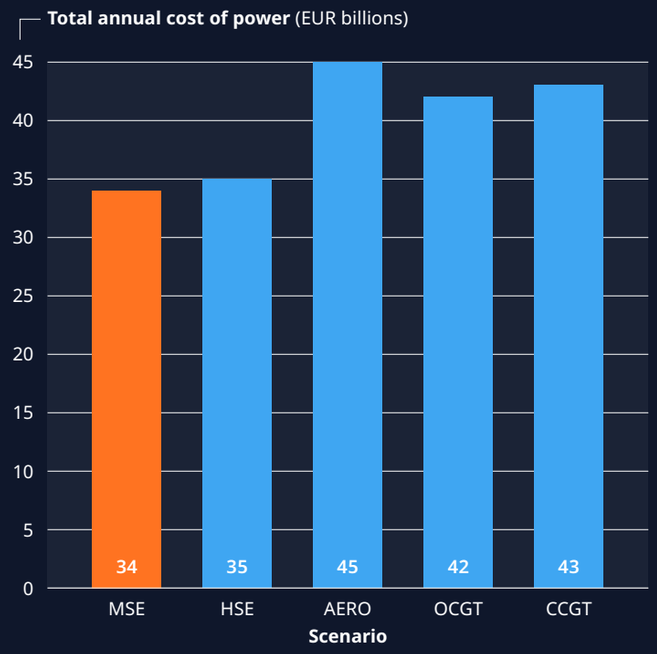
The annual cost of power for onsite-powered data centres includes all costs associated with owning and operating the assets, including annualised capital costs, fixed and variable operations and maintenance (O&M) costs, and fuel costs. The analysis shows that technologies with greater modularity and reliability can reduce required capital and fixed operating costs, while higher fuel efficiency and lower cycling costs can materially lower generation costs over the operating life of the asset.

As shown in Figure 2, these differences become significant when applied across global onsite-powered data centre capacity. **The annual cost difference between medium-speed engines (MSE) and combined-cycle gas turbine (CCGT) scenarios is approximately EUR 9.8 billion**. Over a 20-year operating period, this translates into a cumulative cost difference of approximately EUR 196 billion, illustrating the substantial long-term economic implications of technology selection.

> Technology selection could create approximately EUR 196 billion in cumulative cost differences over a 20-year operating period.

Figure 2

Technology choice creates nearly **€10 billion** in annual power cost differences



Different technologies can have a major impact on long-term operating costs and overall project economics.

Note: Total cost of power includes annualised capital expenditure, fixed and variable operations and maintenance (O&M) costs, fuel costs, start-up costs and applicable carbon taxes. Source: Wärtsilä modelling and analysis (2026).

3.4.

Minimising data centre land use

Land availability is an important constraint in many of the world's largest data centre markets. Data centres are often developed in or near densely populated areas, where suitable sites are already limited and compete with housing, commercial development and other infrastructure. Minimising the land footprint of onsite power solutions is therefore an important consideration.

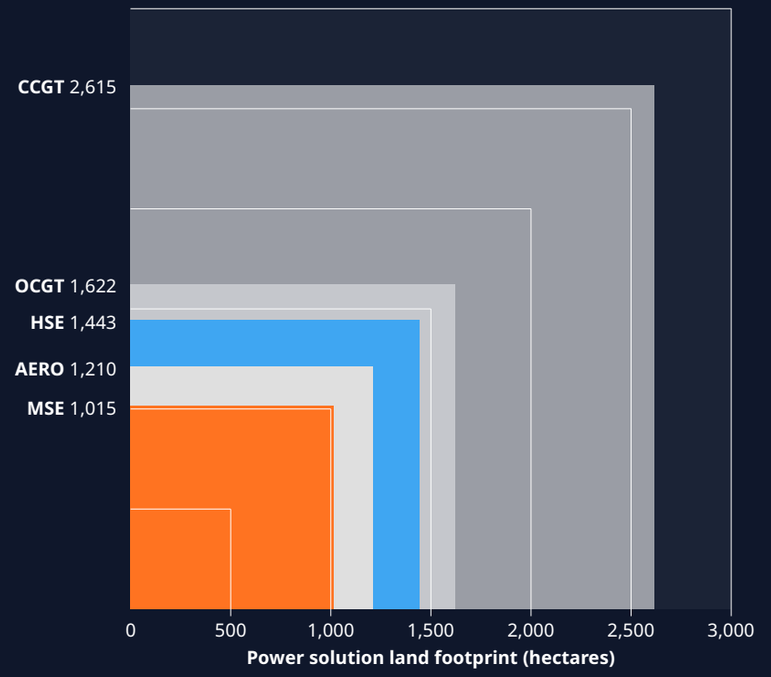
> Combined-cycle gas turbine deployment would require two and a half times more land than medium-speed engines.

As shown in Figure 3, total land requirements vary significantly by power generation technology deployed, ranging from approximately 1,000 hectares for medium-speed engines to over 2,600 hectares for combined-cycle gas turbines (CCGTs). Depending on the technology chosen, the onsite power infrastructure would require an area equivalent to 8–20% of the land required for the data centres.

Some larger generating units can produce more power from a smaller land footprint and require less land per MW of installed capacity. However, this advantage can be more than offset by the additional power generation and energy storage capacity required to meet redundancy and flexibility requirements. As a result, **CCGT deployment would require two and a half times more land than medium-speed engines – approximately 1,600 additional hectares, roughly equivalent to 2,200 football fields.**

Figure 3

Medium-speed engine solution is **over two and a half times more** land-efficient than the most land-intensive option



Land requirements can influence project feasibility, permitting timelines and site selection options.

Note: Project-level land requirement from Power Solutions for off-grid data centres reportt is scaled to global onsite-powered data centre capacity. Single-technology scenarios are illustrative; actual deployment will involve a mix of technologies. The scenarios provide indicative lower and upper bound land use requirements. Land use values reflect the footprint of onsite primary power solutions only, excluding data centre buildings and associated infrastructure. Source: Wärtsilä, 2026.

3.5.

Minimising data centre water use

Water availability is one of the most important considerations for off-grid data centre development, particularly in regions already facing water stress or competing demands from communities, agriculture and industry. As the conversation around data centre water use increases, minimising water consumption is becoming increasingly important to securing and maintaining long-term community acceptance. Developers are therefore under growing pressure to limit the impact of both the data centre and its onsite power solution on local water resources.

As shown in Figure 4, annual water consumption for the onsite power solutions required to supply 55 GW of off-grid data centre capacity varies significantly across the technologies assessed, ranging from less than **0.1 million m³ per year** for medium-speed engines (MSE) to around **314 million m³** for combined-cycle gas turbines (CCGT). This represents a **more than 3,700-fold difference** between the least and most water-intensive power solutions assessed.

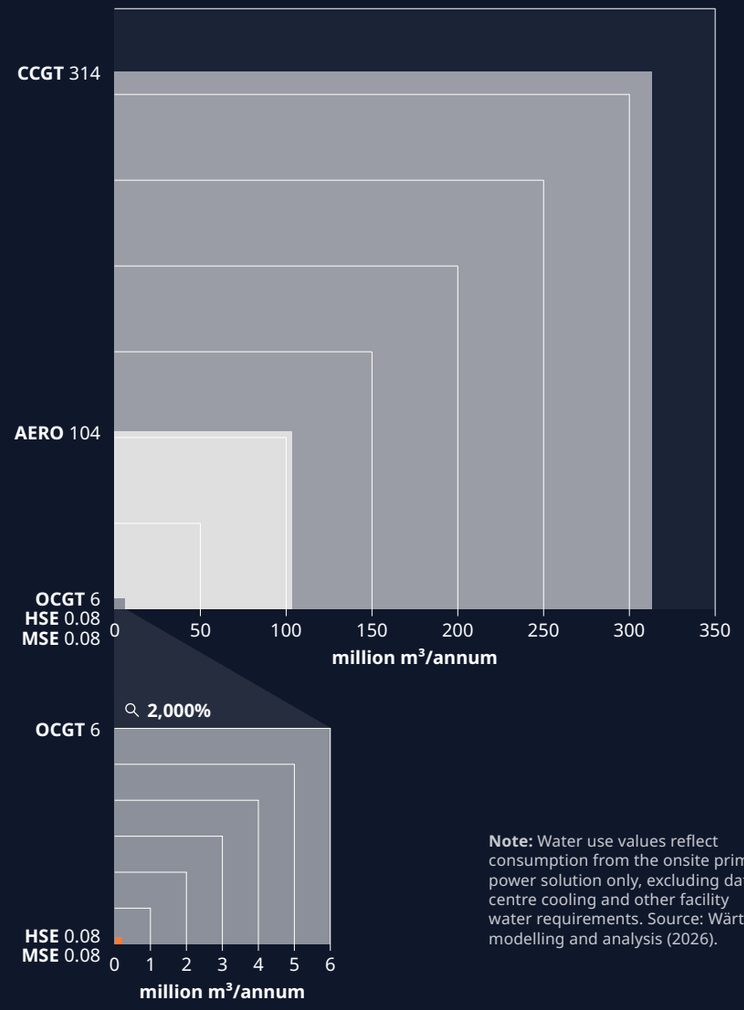
The difference is primarily driven by the cooling requirements across power generation technologies. Some solutions require significant amounts of water for steam-cycle operation, condenser cooling or inlet-air cooling, while others rely predominantly on closed-loop cooling systems. At the scale assessed, the difference in annual water requirements between the solutions with the highest and lowest water-use intensity is equivalent to around 120,000 Olympic-sized swimming pools.

With advances in data centre cooling technology, the water requirements of the data centre itself can be reduced substantially. Direct liquid cooling using electrically driven chillers and air-side economisers can minimise or largely eliminate the need for evaporative cooling, enabling water-use intensities of approximately 0.05–0.4 L/kWh [12]. At the lower end of this range, global onsite-powered data centre water consumption would be approximately 14.4 million m³ annually, equivalent to the annual residential water consumption of a medium-sized city.

> The gap between the most and least water-intensive power solutions equals the annual volume of around 120,000 Olympic-sized swimming pools.

Figure 4

Engine-based solutions deliver by far the **lowest water consumption**



Water consumption varies dramatically depending on technology choice, with implications for permitting, local communities and operating costs.

Note: Water use values reflect consumption from the onsite primary power solution only, excluding data centre cooling and other facility water requirements. Source: Wärtsilä modelling and analysis (2026).

WÄRTSILÄ



Chapter 4

Enabling the path towards a 100% renewable energy future

4.1.

Transitioning beyond off-grid power

As grid connection queues ease, onsite-powered data centres are expected to gradually connect to the electricity grid and other energy networks, such as district heating. As illustrated in Figure 5, this creates a pathway for onsite power solutions to evolve from dedicated primary power sources to helping to balance grids. This allows investment in these power solutions to retain significant value after grid connection.

Once grid-connected, these onsite power solutions can continue to provide backup power to data centres during grid outages while also supporting the wider electricity system. Depending on the technology, location and market design, balancing power could reduce electricity imports from the grid during periods of system stress and, where permitted, provide balancing capacity and ancillary services, such as frequency control and other services that help maintain grid stability. The result is that **infrastructure originally developed to enable rapid off-grid deployment can continue to create value over its operating life, supporting both data centre reliability and the flexibility of the wider energy system.**



Figure 5: Glossary

Onsite power assets can continue **creating value** after grid connection

Types of data centres

On-grid fixed-load data centre

A grid-connected data centre that relies on the electricity grid as its primary source of power. The data centre operates as a fixed load and is not able or willing to reduce its grid electricity demand in response to changing power-system conditions.

On-grid flexible-load data centre

A grid-connected data centre that can flex its electricity demand in response to power-system conditions. This can include reducing demand or using onsite balancing power, helping the data centre maintain operations while providing flexibility and other grid-support services.

Off-grid data centre

A data centre without a connection to the electricity grid. Its electricity requirements are supplied continuously by an onsite primary power solution, either permanently or until a grid connection becomes available.

Assets initially deployed for primary power can later support the grid through balancing and flexibility services.

Types of power solutions

Onsite back-up power solution

An onsite power system designed primarily to provide emergency electricity supply when grid power is unavailable.

Onsite balancing power solution

An onsite power solution at an on-grid flexible-load data centre that provides back-up power during grid outages and flexible power during normal grid-connected operation, enabling the data centre to reduce its grid demand and provide grid balancing or other system-support services.

Onsite primary power solution

An onsite power solution designed to continuously supply the electricity requirements of an off-grid data centre.

Offsite power solution

A power generation facility located away from the data centre and connected to the electricity grid. The data centre may procure its electricity through the grid via a dedicated or contracted power plant, renewable energy project or another generation asset, including through a power purchase agreement (PPA).

4.2.

Unlocking a new source of grid flexibility

Data centres' onsite power solutions could become a significant source of flexibility for the wider electricity system, reducing the need for dedicated balancing capacity and lowering the cost of system balancing, while integrating growing volumes of renewable generation.

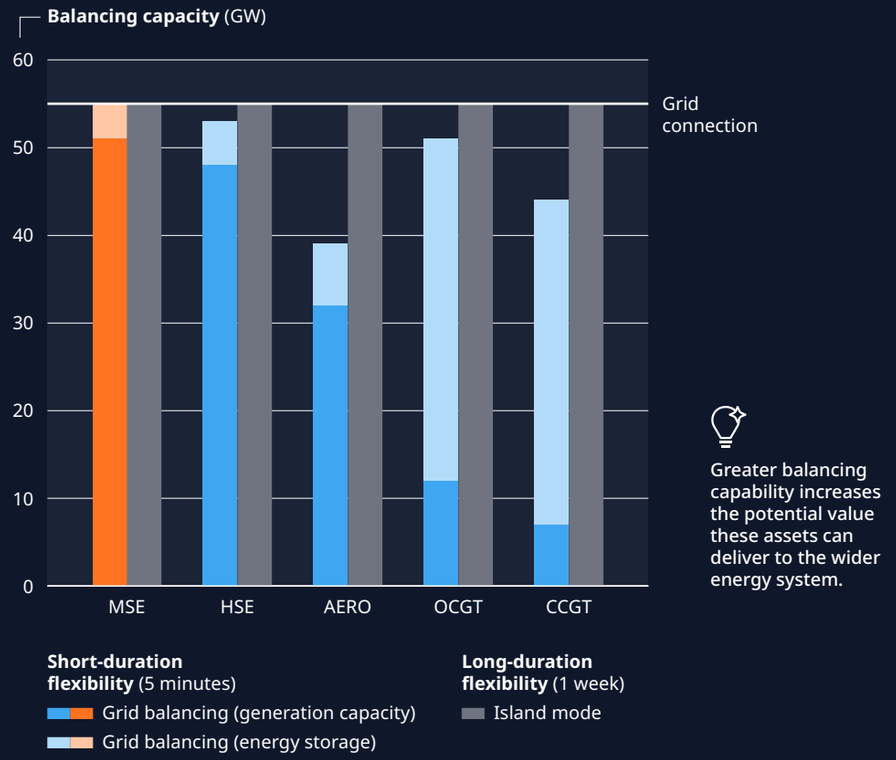
> Data centres could evolve from power consumers into providers of grid flexibility.

As illustrated in Figure 6, this balancing could be provided through two distinct operating modes. Short-duration balancing could be delivered by the onsite power solution, exporting power to the grid to help balance short-term renewable variation. Long-duration flexibility could be provided by data centres temporarily operating independently from the grid, known as island mode, during periods of sustained grid stress, such as prolonged periods of low renewable output or high system demand. In this mode, the data centre would disconnect from the grid and meet its electricity needs using onsite power instead.

The choice of onsite power solution materially affects both the level of short-duration balancing capacity available to the grid and the volume of renewable capacity that this flexibility can support. Of all the power solutions that could be deployed to support 55 GW of onsite-powered data centres globally, medium-speed engines would provide the greatest grid-balancing capability due to their high availability and operational flexibility.

Figure 6

Medium-speed engines provide the **greatest grid-balancing capability**



Note: Short-duration flexibility reflects the five-minute ramp capability of installed onsite capacity, capped by a grid connection equal to peak data centre demand (55 GW). Long-duration flexibility assumes the data centre can transition to island mode during prolonged periods of grid stress. Only primary generating capacity and ramp reserve energy storage capacity is assumed to be available for grid balancing. Power generation and energy storage capacity reserved for maintenance and instantaneous reserve are excluded to ensure that participation in grid balancing does not reduce data centre power reliability in the event of a sudden and prolonged loss of grid connection. Source: Wärtsilä 2026.

WÄRTSILÄ

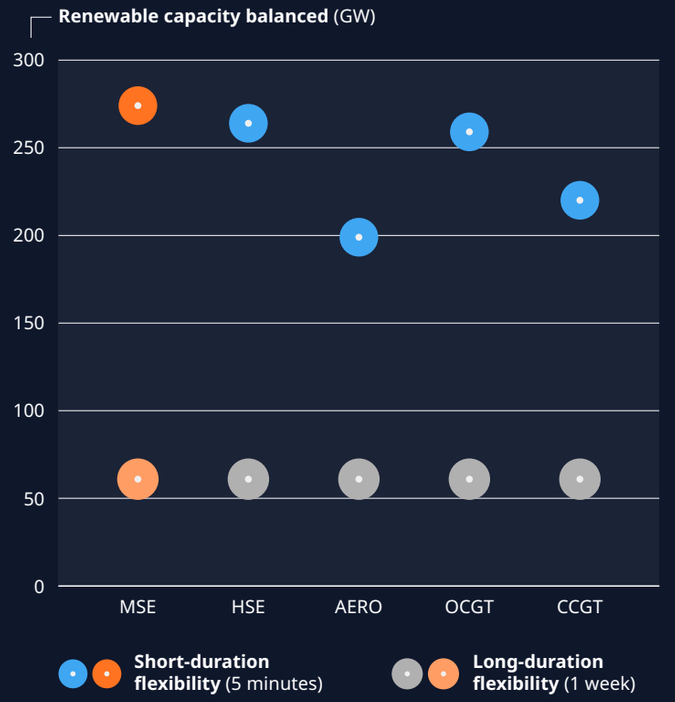
As shown in Figure 7, 55 GW of medium-speed engine based power solutions could provide sufficient short-duration balancing to manage the five-minute variability of approximately 270 GW of renewable capacity. Across the technology options assessed, broadly similar levels of long-duration flexibility are available by allowing data centres to disconnect from the grid and operate in island mode during prolonged periods of grid stress. Under this approach, the same onsite power solutions could provide long-duration flexibility equivalent to approximately 61 GW of renewable capacity.

> 55 GW of onsite data centre flexibility could avoid up to EUR 10 billion per year in equivalent system capacity costs.

The potential economic value of using onsite data centre power assets to provide balancing capacity is substantial. At annualised capital and fixed operating costs of around EUR 185/kW-year for new balancing capacity, unlocking **55 GW of system flexibility could avoid up to EUR 10 billion per year in equivalent system capacity costs**, or approximately EUR 200 billion over 20 years. This would mean, in effect, that data centres help finance the balancing capacity that the wider electricity system needs to support renewable integration, helping to partly overcome the "missing money" challenge faced by new balancing power plants. By reducing the need for additional dedicated capacity, this could lower overall electricity system costs and enable higher shares of renewable energy to be integrated into the grid.

Figure 7

55 GW of onsite power could provide the short-term balancing needs of **270GW of renewable capacity**



Flexible power assets can help integrate significantly more renewable energy into the grid.

Note: Short-duration balancing assumes renewable generation can vary by up to 20% over five minutes. Long-duration balancing assumes renewables provide a 10% effective load-carrying capability (ELCC). Source: Wärtsilä 2026

WÄRTSILÄ

4.3.

Reimagining data centre backup power

Most grid-connected data centres install onsite backup power solutions to maintain operations during grid outages. However, this capacity is typically provided by high-speed diesel generators designed almost exclusively for emergency use. Their low efficiency, high operating costs and local air emissions make them poorly suited to frequent grid balancing operation. As a result, there is substantial onsite generation capacity deployed globally today that remains largely idle even as electricity systems face a growing need for capacity and flexibility.

With around 180 GW of new data centre capacity expected globally over the coming decade, creating incentives for grid-connected facilities to deploy higher efficiency, flexible power generation, such as engines and energy storage that are designed for both backup and balancing power, could unlock a major new source of energy system flexibility and capacity. If such solutions were adopted across the data centre sector, the potential flexibility and capacity resource could theoretically reach up to 180 GW, equivalent to around 42% of the gas-based power generation expected to be added globally by 2035 under the IEA's Stated Policies Scenario [2]. In this way, data centres could help finance a large part of the balancing infrastructure required to integrate renewable generation while simultaneously strengthening their own resilience.

4.4.

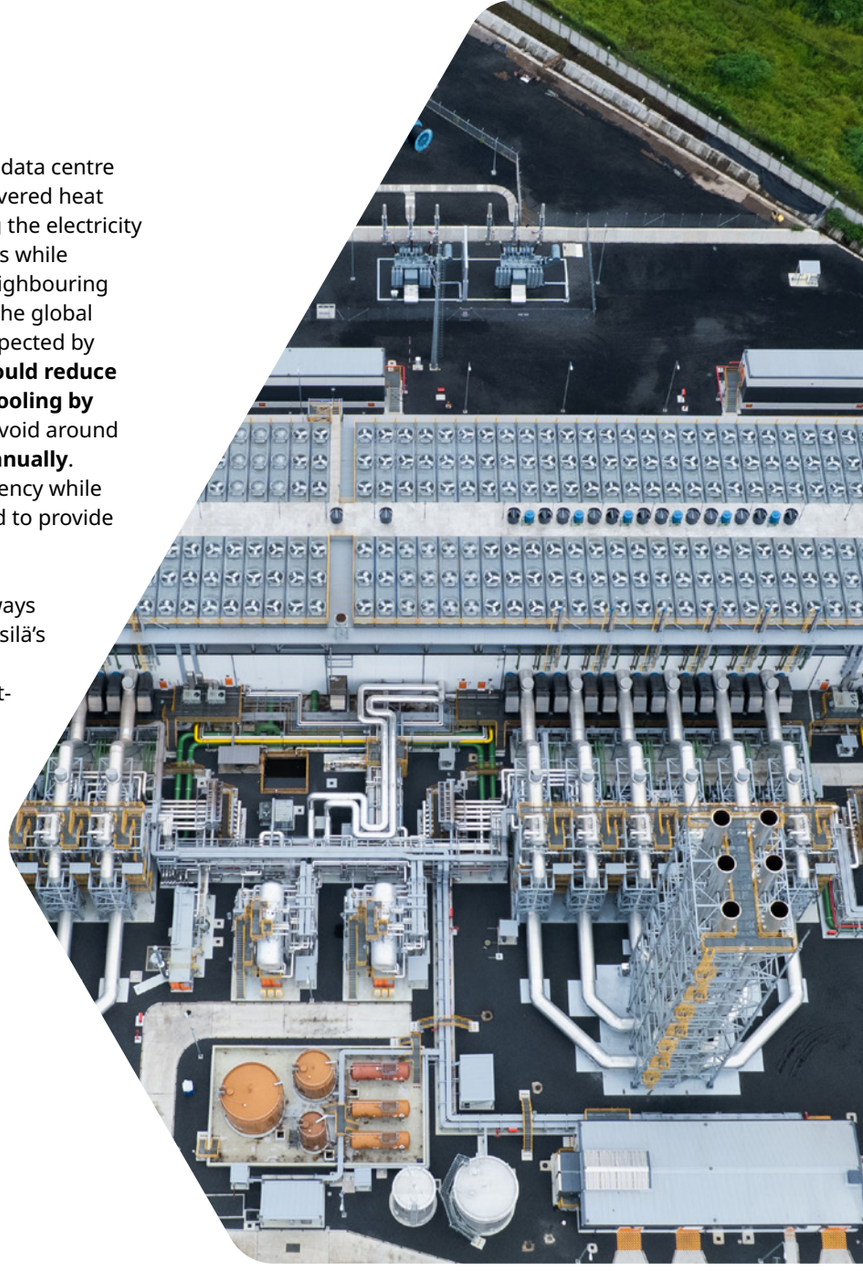
Turning waste heat into value

Waste heat recovery can increase the economic value of onsite primary power solutions serving off-grid data centres. Gas-based power generation produces substantial quantities of medium- to high-grade heat that would otherwise be rejected to the environment. Where suitable heat demand exists nearby, this recovered energy can be used for industrial processes, district heating networks, commercial buildings and other applications, improving overall energy utilisation and creating an additional source of value from the onsite power solution. The opportunity depends primarily on the temperature and quantity of recoverable heat, proximity to suitable users and the availability of heat-recovery and distribution infrastructure.

> Data centres could provide up to 180 GW of flexible capacity to the power system.

Waste heat can also be used within the data centre itself. Through absorption chillers, recovered heat can be converted into cooling, reducing the electricity required by conventional electric chillers while also potentially supplying cooling to neighbouring commercial or industrial users. Across the global onsite-powered data centre capacity expected by 2035, **deploying absorption chillers could reduce electricity consumption for electric cooling by approximately 39 TWh per year** and avoid around **19 million tonnes of CO_2 emissions annually**. This could improve overall system efficiency while reducing the electrical capacity required to provide data centre cooling.

However, waste heat recovery is not always cost-effective. As demonstrated in Wärtsilä's *Power solutions for off-grid data centres* study, the additional capital cost of heat-recovery equipment and absorption chillers can outweigh the savings from reduced electricity consumption in some data centres. Where policymakers or system operators want to increase the use of otherwise wasted heat, targeted incentives may be needed to strengthen the investment case.



Chapter 5

From servers to system resources



5.1. The limits of flexible compute

Demand-side flexibility, where data centres adjust their electricity consumption in response to grid conditions, is increasingly being explored as a way to support power systems, particularly during periods of high system demand, constrained grid capacity or high electricity prices. From the grid's perspective, the effect is similar whether a data centre reduces its load or offsets part of its electricity demand by operating onsite generation: both reduce net demand on the electricity system.

However, several studies have highlighted the practical limitations of relying on load flexibility, suggesting that its contribution to overall grid flexibility may be limited [4,5]. AI inference is expected to become the dominant workload type by 2030 and generally requires low latency, high availability and continuous operation [4,5]. AI training may offer greater flexibility, but the significant capital invested in computing infrastructure creates a strong incentive to maximise server utilisation. Flexibility is further constrained in colocation facilities, where operators typically do not control tenant workloads. IT-load flexibility is therefore likely to be concentrated in selected non-critical workloads and limited temporal or geographic workload shifting [4].

5.2. Harnessing waste heat from data centres

Computing infrastructure produces significant quantities of heat that would otherwise be released to the environment, creating an opportunity to recover and reuse it where suitable heat demand exists. The potential value depends primarily on proximity of heat users, the temperature and quality of the recoverable heat, and the availability of supporting infrastructure.

Where conditions are favourable, recovered heat can be supplied to district heating networks, industrial facilities and large commercial users. The growing adoption of direct liquid cooling for high-density AI workloads can further improve this opportunity by increasing the quality of recoverable server heat, making it more suitable for productive reuse.

The opportunity is particularly significant in Europe, where extensive district heating networks provide an established route to market. The IEA estimates that data centres in Europe could potentially supply around 300 TWh of heat, equivalent to approximately 10% of European space-heating demand [1]. At an assumed heat tariff of EUR 50/MWh_{TH}, this represents potential gross value of around EUR 15 billion per year. After accounting for the electricity required by heat pumps to raise recovered heat to useful temperatures, our analysis indicates **potential net economic savings of approximately EUR 9 billion annually, or over EUR 180 billion over a 20-year period.**

Chapter 6

Pathways for net zero data centres

> Data centres could not only support the deployment of additional renewable capacity, but also help integrate and balance the renewable electricity they support.

Onsite power solutions can accelerate speed-to-power without locking data centres into a long-term, high-emissions operating model. Over time, they can provide a bridge to grid connection, enabling data centre capacity to come online sooner before transitioning to a balancing and back-up power role as additional grid capacity becomes available.

As illustrated in Figure 8, the study outlines a transition pathway in which onsite power solutions progressively shift from primary power supply to providing system balancing. As grid access expands, grid electricity increasingly displaces onsite generation. At the same time, long-term renewable power offtake by data centres can support investment in new renewable generation, increasing the supply of renewable electricity available to the power system and reducing the carbon intensity of both data centre electricity consumption and the wider power sector.

Rather than becoming stranded once grid connection is established, onsite power solutions can transition into balancing power solutions. They can continue to supply the data centre during grid outages and prolonged periods of system stress, while also providing short-duration balancing and ancillary services by exporting power to the grid. Data centres could therefore play a dual role: supporting investment in new renewable capacity through long-term electricity procurement while also providing the flexible capacity needed to balance variable renewable generation across the wider power system.

The fuel used by balancing solutions over the life of the asset will depend on regional economics, policy and fuel availability. Natural gas may remain suitable for low-running-hour balancing and backup in some markets, while higher gas prices, stronger carbon pricing or supportive policy frameworks could encourage a gradual transition towards low-carbon or sustainable fuels such as biomethane, renewable methane, hydrogen or ammonia as supply chains mature.

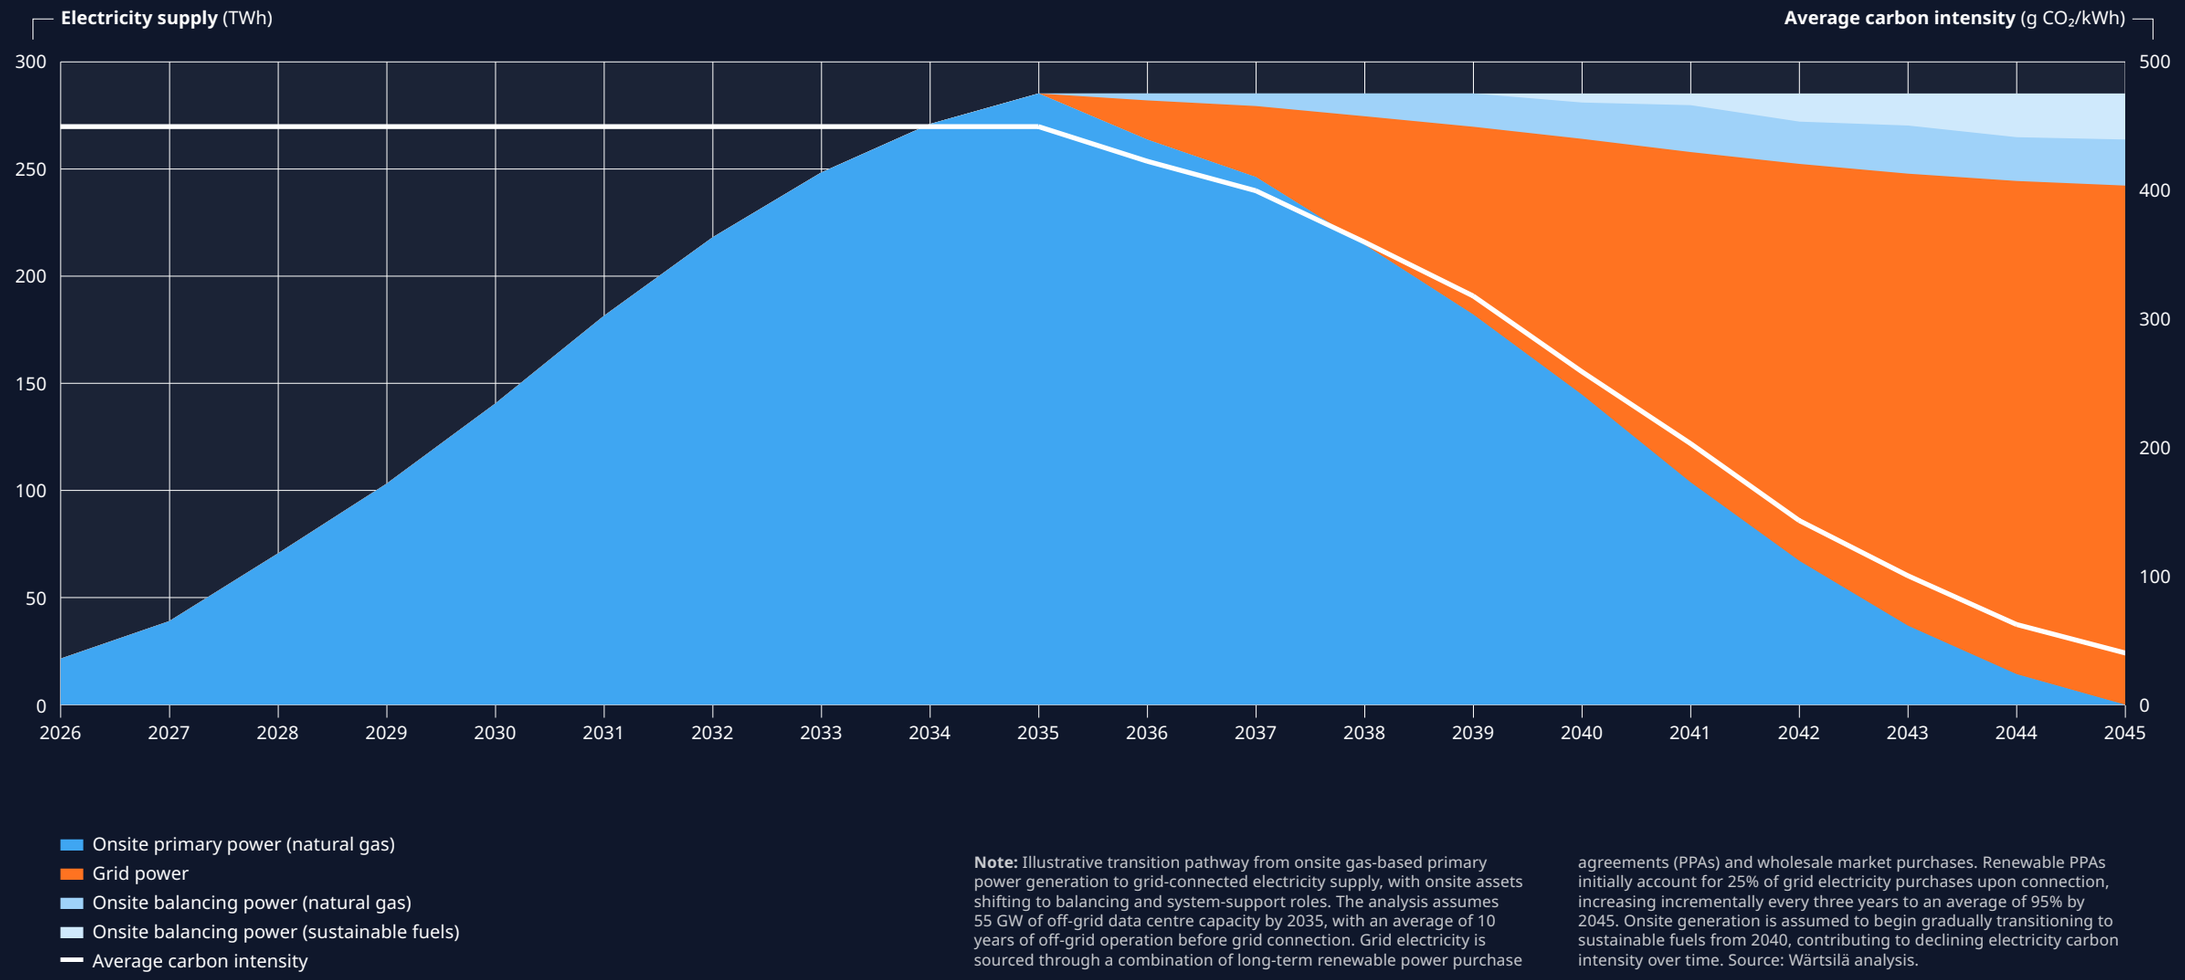
Under the pathway illustrated, renewable electricity progressively becomes the principal source of energy, while onsite power solutions shift from continuous primary power supply to lower-running-hour balancing and backup operation. Even if only around 50% of balancing capacity uses sustainable fuels by 2045, the combination of substantially lower onsite generation and a much higher share of renewable electricity from the grid could reduce the average carbon intensity of data centre electricity to around 40 gCO_2/kWh.

The feasibility and system value of this pathway will ultimately depend on the technology selected for the initial off-grid stage. Technologies combining high efficiency, rapid ramping, high availability and operational flexibility can transition more effectively from continuous primary power to balancing services. Less flexible solutions may provide fewer system services and contribute to higher renewable curtailment, limiting their potential value to the wider electricity system.

Figure 8

Following grid connection, combining renewable energy sourcing with balancing power solutions can provide **a pathway towards net-zero emissions** for data centres

Operational emissions depend not only on technology choice, but also on how assets are integrated into the broader energy system.



Note: Illustrative transition pathway from onsite gas-based primary power generation to grid-connected electricity supply, with onsite assets shifting to balancing and system-support roles. The analysis assumes 55 GW of off-grid data centre capacity by 2035, with an average of 10 years of off-grid operation before grid connection. Grid electricity is sourced through a combination of long-term renewable power purchase agreements (PPAs) and wholesale market purchases. Renewable PPAs initially account for 25% of grid electricity purchases upon connection, increasing incrementally every three years to an average of 95% by 2045. Onsite generation is assumed to begin gradually transitioning to sustainable fuels from 2040, contributing to declining electricity carbon intensity over time. Source: Wärtsilä analysis.

WÄRTSILÄ

Chapter 7

Conclusion

The rapid growth of off-grid data centres makes technology selection increasingly important not only at the project level – but also for the wider energy system. The choice of onsite power solution directly influences the generation and energy storage capacity required, capital investment, operating costs, land use, water consumption and exposure to equipment supply constraints. More modular, reliable and flexible technologies can reduce infrastructure requirements while improving the speed, resilience and economics of data centre deployment.

The opportunity extends well beyond the initial off-grid phase. As grid connections become available, onsite power assets can continue to deliver value rather than becoming stranded infrastructure. They can instead transition from primary power solutions into balancing power solutions, maintaining data centre reliability while reducing grid imports during periods of system stress and, where market arrangements allow, providing balancing, reserve and capacity services.

> With the right technology choices and policy framework, onsite power can evolve from a speed-to-power solution into a long-term system asset.

Leveraging this existing infrastructure could reduce the need for separate investment in new flexible capacity, lower the cost of integrating renewable generation and help address shortages of balancing capacity that can constrain further renewable deployment. In this way, investment to accelerate data centre development can also help finance the flexibility required by the wider electricity system.

Technology choices made today can therefore have broader system implications tomorrow. Well-designed regulation, market incentives and grid arrangements that recognise and reward flexibility in onsite generation could align data centre investment with future power system needs, creating synergies between faster data centre deployment, renewable integration and lower system costs.

With the right technology choices and policy framework, onsite power can evolve from a speed-to-power solution into a long-term system asset, supporting more resilient data centres, a more flexible electricity system and the integration of significantly greater volumes of renewable electricity.

References

1 International Energy Agency, Energy and AI, 2025

2 International Energy Agency, World Energy Outlook 2025, 2025

3 International Energy Agency, Electricity 2026, 2026

4 ENTSO-E, Data Centres and the power system: expected trends, challenges and opportunities, 2026

5 McKinsey & Company, The Future of AI Workloads, 2026

6 International Energy Agency, Key Questions on Energy and AI, 2026

7 Berkeley Lab, Queued Up: 2025 Edition, Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2024, 2025

8 BloombergNEF, Data Centers Eye 114GW of On-Site Gas Amid Grid Delays, 2026

9 S&P Global, Pipeline operators strike deals as data centers turn to colocated generation

10 BloombergNEF, Six Things to Know About BNEF's New US Data Center Capacity Outlook

11 Cleanview, Bypassing the Grid: How Data Center Developers Are Building Their Own Power Plants

12 Berkeley Lab, 2024 United States Data Center Energy Usage Report, 2024

13 BloombergNEF, US Data Center Outlook 1H 2026, 2026

WÄRTSILÄ


WÄRTSILÄ