Applying Natural Gas Engine Generators to Hyperscale Data Centers

White Paper 286

Version 1

by Wedian Youssef (Schneider Electric), Adam Rajewski (Wärtsilä), Malik Megdiche (Schneider Electric) and Juha Kerttula (Wärtsilä)

Executive summary

Power capacity shortages, a desire to be independent from the grid, and increasing pressure to reduce emissions, are all drivers for having self-generating power on site. Providing a reliable, environmentally-friendly power supply, however, can be challenging. We will show that natural gas-fired, on-site generation power plants for either backup or primary power supply functions, can be an attractive alternative to diesel-based power plants. We will explain the technical requirements and design adaptations needed for hyper-scale data centre applications and will compare our proposal to traditional diesel generation power plants. Finally, we explain how to extend the benefits of gas-engine technology and turn the generation plant into a revenue-generating asset.

Introduction

The gas engine generation plant described in this paper is based on current state-of-the-art, medium-speed, gas-engine technology. During recent years, this technology has become more common and applied to increasing numbers of power generation applications. The reciprocating engines used represent the most efficient simple cycle power generation technology available today, while at the same time they can perform extremely fast start-ups and handle sudden load changes. This makes them suitable for industrial and utility-scale applications, ranging from commercial power peaker plants to various on-site and off-grid power generation solutions – including mission-critical data centres. The largest engine power plants currently being supplied have an installed capacity of approximately 600 MW. Given the typical output of medium-speed gas-engines, they are generally more suited for larger data centres with a minimum continuous load beyond 10 MW, ideally around 20 MW.

In order to propose an effective data centre design, we must first consider customer needs such as the facility's availability requirements, the electrical equipment to be used, and the chosen electrical distribution architecture. A variety of architectures could be proposed for large scale data centres to meet availability and reliability requirements. These various architectures rely mainly on the utility as the principle source feeder with backup generation from a diesel generator power plant, either at medium voltage (MV) or low voltage (LV) depending on the architecture choices and the size of the site. This paper will consider N, N+1, and 2N type architectures. Depending on the uptime needed, redundant architectures using "2N" or "N+1" redundancy levels are most common. A typical solution for a large-scale data centre is presented in **Figure 1** and is shown alongside an architecture using natural gas generators.

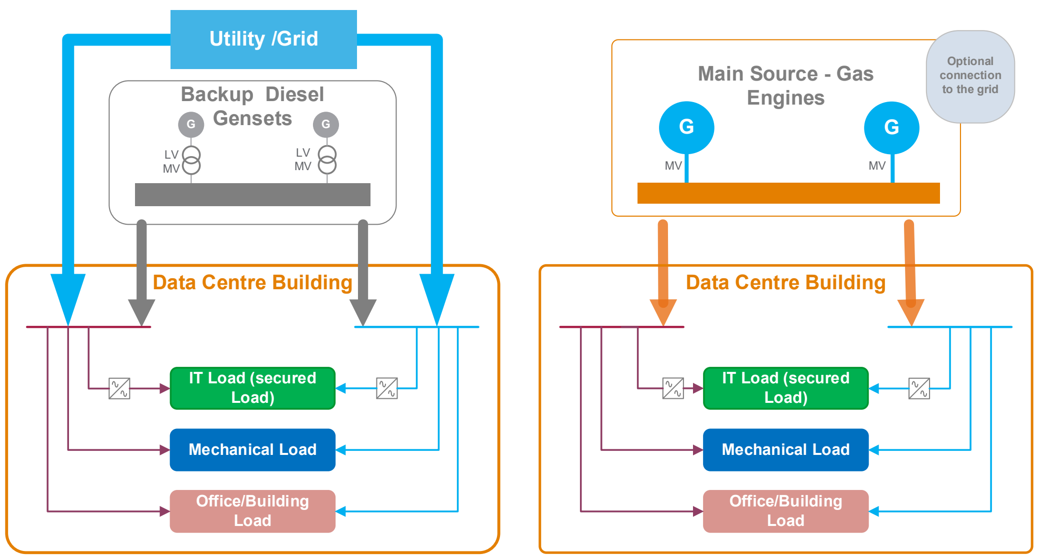

Figure 1
Block diagram showing the difference between a traditional diesel-powered solution and a gas-powered solution without connection to the grid

A shift in architecture from the utility power supply to on-site, gas-fired generation requires rethinking the design of the data centre in order to be confident with the proposed solution. Data centre loads are not always constant in power draw and require a continuous supply of power to maintain availability. An ISO Continuous Operation Power (COP) -rated generator might meet the unlimited run-time requirement for Tier III and IV data centres. However, it may not be the best fit for a low load profile. Given this consideration, generator manufacturers have come up with generator designs that fully meet data centre application requirements. And to fulfill this requirement, we must go beyond the engine design itself and take into consideration multiple factors, including the availability of fuel, the electrical characteristics (e.g., transient voltages, frequency deviation), recovery time, and emission standards.

Natural gas generation plant characteristics

Reduced environmental footprint

The most fundamental differentiating characteristic of gas engine technology from diesel engine technology is its reduced environmental footprint. This results from the combination of an inherently "cleaner" fuel, and a more efficient engine combustion cycle. This means a low carbon footprint of just above 400 g/kWh. This is not only lower than for other on-site power generation technologies, such as diesel engines, but it is also lower than the average value for grid power in many areas of the world, as shown in **Figure 2**. Furthermore, emissions of local pollutants are considerably less than those from diesel-powered generators (**Figure 3**). This makes it easier to obtain environmental permits and social acceptance for projects with on-site generation based on natural gas, than it would be for diesel generators.

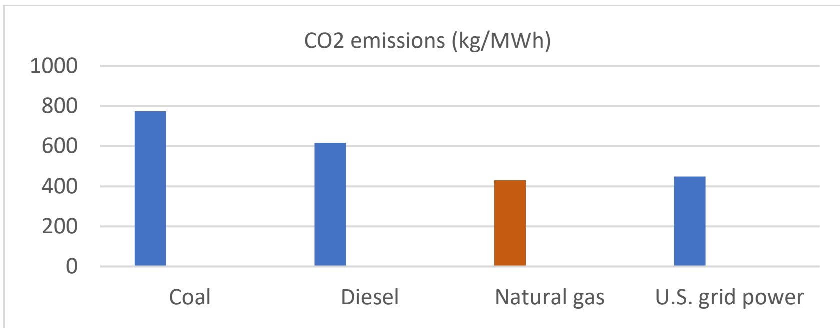

Figure 2
The carbon footprint for different sources of electricity. Note that generating power with an on-site natural gas power plant may have a smaller footprint than grid electricity (based on Wartsila calculations, data for U.S. grid by U.S. Energy Information Administration for 2018).

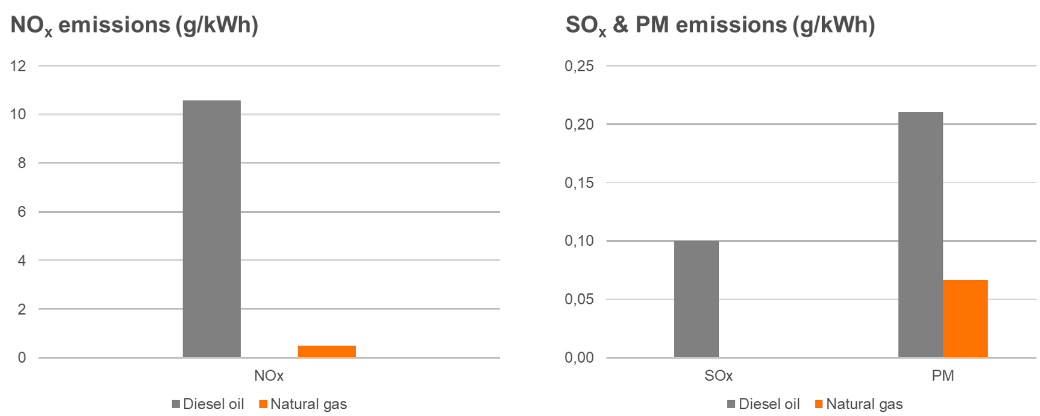

Figure 3
Gas is a much cleaner fuel than diesel oil. Power generation from natural gas results in far fewer emissions of all pollutants. Therefore, gas power plants may be freely operated and not limited to emergency situations. Note that the values provided are achievable with engine tuning alone; NOx emissions may be further reduced using selective catalytic reduction (SCR) technology.

Gas engine attributes important for large data centre applications

Gas engines used as prime movers in the solution described in this paper are spark-ignited engines fueled with natural gas, operating at 750 revolutions per minute, with a mechanical output of 10 MW each, achieving a high generation efficiency of 46-47%. These engine-generator sets can be started, synchronized, and can reach nominal speed within 15 seconds. After synchronization, supporting the site's full load takes an additional 35 seconds. This can be reached either through linear loading or by taking the load in steps, each of which may reach approximately 20% of the nominal output (**Figure 4**).

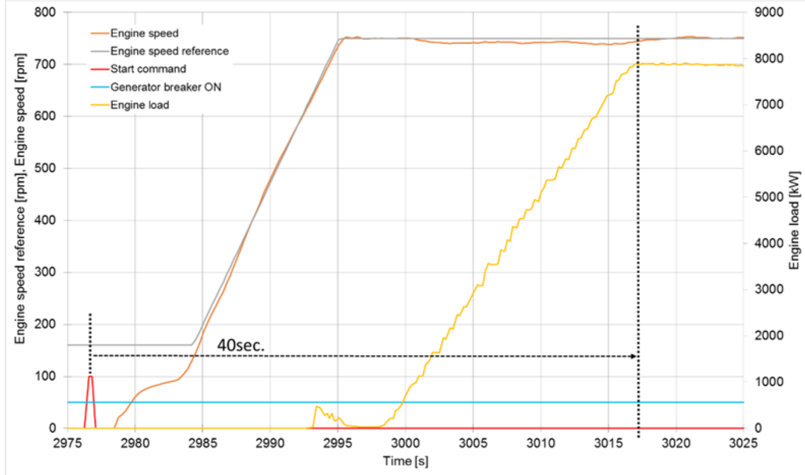

Figure 4
Start-up curves of a modern medium-speed gas engine. These are direct screenshots from engine control systems and were made during tests of a medium-speed gas engine's rapid start-up while operating in island mode. The start-up duration is defined as the time taken from the start command until full output is achieved.

Because each engine in a power plant is functionally independent from the others, the system can provide its full power within a minute of the start command, regardless of the total output. This makes it well-suited for emergency applications where critical loads are protected by UPS units. However, unlike typical emergency diesel generators, these gas engines are designed for continuous base load operation. This means that a power plant built using gas engines may be operated in almost any mode – from pure standby to continuous operation. Properly designed plant systems ensure concurrent maintainability and enable the creation of a completely "off-grid" system supplying power 24/7 without any interruptions. Functionally, a multi-engine power plant is essentially a set of individual single-engine power plants. The special features of the gas engine technology and power plants dedicated for data centre applications are discussed in more detail in the following sections.

Gas engine technology

The gas engine is the heart of the solution discussed in this paper. Generally, two types of gas engines, high-speed and medium-speed, are used in modern power engineering. High-speed technology, with speeds in excess of 1000 rpm, is used in smaller scale engines of up to 4 MW per unit, whereas medium speed technology (typically 750 rpm for 50 Hz and 720 rpm for 60 Hz) is used for sizes above that. The differences between these technologies are as follows:

- Medium-speed engines are somewhat larger and sturdier – they need a larger volume to achieve the same power output. This makes them also somewhat more expensive to construct and install.
- Medium-speed engines are less expensive to maintain since their components have a slower rate of wear. This means a lower frequency of major overhauls, and a fewer number of parts that need to be reconditioned or replaced.
- High-speed engines typically feature a single-point gas admission system. Gas is fed into an air inlet duct before the turbocharger. This system, while simple and inexpensive, also has significant disadvantages. It does not permit individual cylinder control and it causes delay in the engine control system, making the engine unable to start quickly and less able to handle rapid load changes. In medium-speed engines, gas is typically supplied individually to every cylinder's inlet. While this requires a somewhat higher gas supply pressure (around 6-7 bar or 85-100 psi), it brings significant benefits. Most importantly, this allows optimising the gas dosage to match the conditions of each single cylinder and every single cycle. This enables operating closer to the knocking and misfiring limits, which ultimately increases engine efficiency. Another important result is a much better control response, since the fuel doses can be adjusted

virtually without delay, which makes this type of engine better at taking rapid load steps. They also react better to varying gas conditions.

From a combustion cycle point of view, medium-speed gas engines can be further divided into spark-ignited gas engines with a so-called Otto-cycle (i.e. the same as in petrol-fuelled car engines) and dual-fuel engines utilising a Diesel cycle. **Engines with spark ignition**, which only burn gaseous fuels, represent the most efficient option. Such engines are characterised by having the highest efficiency and lowest emission levels. Fuel storage may be provided in the form of liquefied natural gas tanks, as discussed further below.

The alternative solution are **dual-fuel engines**, which can operate either on liquid or gaseous fuels. In liquid fuel mode, the engines operate essentially as standard diesel engines, with just as much flexibility (the ability to start fast, handle load changes, etc.). In gaseous mode, they operate in a cleaner and more efficient way, although they do require a very small and continuous dose of liquid fuel for ignition, as the engines are not equipped with spark plugs. Dual-fuel engines may switch between fuels when operating at any load, so should the gas infrastructure unexpectedly fail, dual-fuel engines will simply automatically switch to locally stored diesel fuel. This ultimate flexibility comes at the cost of having a slightly lower efficiency than that of “pure” diesel or gas engines, and a somewhat higher level of emissions, which must be reduced with proper after-treatment. With liquid fuels there is also a risk of spillage, which does not exist with natural gas as it will simply evaporate in the event of a leak.

Maintainability

Unlike some smaller, high-speed emergency diesel generators, the medium-speed gas engines used in the commercial power industry are designed to be maintained on site. All maintainable components, such as the cylinder heads, cylinder liners, pistons, connecting rods etc., may be removed using an on-site overhead crane and transported individually to refurbishment workshops, while the heaviest components – the engine block, generator and crankshaft – remain at the site for the lifetime of the plant, which may easily exceed 25 years of base load operation.

In gas engine power plants with medium-speed engines, engine overhauls are required after no less than 16,000 running hours, and the most extensive overhaul likely won’t be needed until between 64,000 and 96,000 running hours (i.e. after 8 to 12 years of continuous operation). This means that for peaker power plants the largest overhaul may not ever be needed during the project’s economic lifetime. It should be noted that the number of starts and stops made by the engine has no impact on its maintenance schedule, unlike in some other power generation technologies, where a single start may be counted as being equivalent to multiple running hours due to the thermal stress added on startup. This maintainability is an inherent feature of the gas engine technology, designed for cyclic operation.

Storing gas on site for increased availability

Spark-ignited, gas-fired engines in most cases operate on natural gas supplied by a regional or national gas pipeline network. Gas networks are known for their very high reliability, but should there be doubts regarding the network, fuel can also be stored on site using liquefied natural gas (LNG) technologies that allow for more compact storage volumes. The storage itself is in well-insulated cryogenic tanks. With small scale installations they are mostly made of steel. This technology is well proven and is very safe. It is currently often used on ships (including cruise vessels carrying commercial passengers), where it allows far more favorable emission characteristics than with previously used diesel oils. The same engine technology can utilise gaseous fuels of biological origin, such as biogas or biomethane. These fuels can also be stored in liquefied form.

Storing natural gas in liquefied form allows for a very high storage density, resulting in a smaller required storage volume. For example, in order to secure 12 hours of operation with a 10 MW engine, a volume of less than 50 m³ (~1700 ft³) would be sufficient. Due to the physical constraint of retaining its liquid form, the fuel needs to be kept at a low temperature (below minus 162 °C or minus 260 °F). This is ensured by using well-insulated cryogenic tanks. As already mentioned, tanks of this type are now becoming common in the shipping industry and are delivered as modular structures.

Liquefied gas is now becoming a widely available commodity that can be freely purchased in many industrialised countries. For multi megawatt-sized facilities, it is typically transported by road (sometimes also by barge). An alternative solution could be provided by the on-site liquefaction of pipeline gas or locally generated biogas.

Technical description of the off-grid solution

Moving from a standard grid-connected solution, with diesel-powered generators used as back up, to an on-site generation solution requires multiple considerations regarding the design of the power plant. Depending on the required reliability, the plant infrastructure must be adapted to meet that targeted requirement. Acknowledging that the main concern is electrical availability, the common parts (electrical and non-electrical) contributing to the generation play a critical role in the reliability and their design should be considered carefully. **The following sub-sections describe the off-grid solution developed from a joint study between Wartsila and Schneider Electric that is optimised for hyper-scale data centre applications.**

A gas power plant architecture overview

Figure 5 provides a high-level illustration of the power plant components. The engine itself and its auxiliaries are considered as one unit, and we can clearly highlight some common systems, including the gas piping, LNG storage, lubrication system, fire detection system, electrical distribution, and the control systems. These individual systems must be analysed and designed as a complete integrated system in order to avoid any single points of failure. The purpose of the analysis is to detect any common failures and create a workaround to avoid an outage by adding redundant components.

In the generation plant analysed and developed for this paper, some of the auxiliary systems, such as the starting air unit and controller unit, are dedicated to the engines, and the common sub-systems have been divided into critical and non-critical. The sub-systems must be fully redundant and designed without a single point of failure. For example, the gas pipeline main feeders are redundant and physically separated from each other. Also, on-site gas storage, usually using LNG, is required. The storage volume depends on the gas availability and how easy it is to access the supply. For this study, we set the requirement at 48 hours assuming a full load capacity.

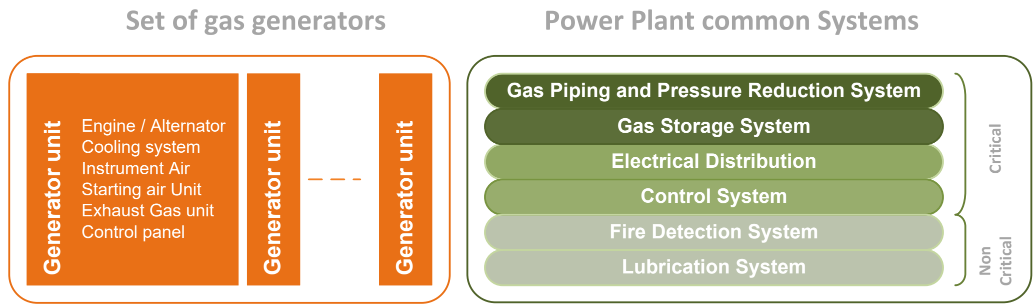

Figure 5
A high-level presentation of the power plant components

In this study to develop an electrical design, we targeted what is considered an acceptable level of availability, which in the hyper-scale data centre market is 99,999%. This is equivalent to five minutes of downtime per year. For data centres, the unexpected event to avoid is a loss of power to the IT servers for more than 20 ms. Given that the servers are backed up by UPSs, the generator plant has about 5 minutes to come online and support the load, if paths A and B are lost, given typical UPS battery runtimes. In the context of on-site generation, the generation set becomes critical, therefore. And so the medium voltage portion of the electrical network design becomes a key focus area in order to achieve the required availability for the overall system.

Electrical distribution architecture

The design challenge was to create a power plant with multiple 10 MW gas engine units with MV alternators that are rated up to 15 kV. To minimise the level of current, including short-circuit currents, the first idea might be to use a step-up transformer for each generator. This, however, is not the optimal choice because it would add significant capital and operating expense. The better option, we found, was to target a single MV level without step-up transformers by using 15 kV for the entire data centre MV distribution system. This allows the use of the more cost-effective 17.5 kV MV equipment range with a short-circuit withstand rating of 31.5 kA maximum, and with a rated current of 3150 A maximum.

Short-circuit constraints

The design goal was to put as many generators in parallel as possible and still be able to use circuit breakers in the 17.5 kV range with 31.5 kA rms breaking capacity, along with 17.5 kV-rated switchboards with a 31.5 kA rms short-circuit withstand (and 79 kA peak). The main values to consider are the rated peak withstand current (that represents the electrodynamic constraint), the rated short-circuit breaking current (that represents the circuit breaker's ability to break the current) and the rated short-time withstand current during 1 second (that represents the thermal constraint during a short-circuit). The above design goal is intended to optimise the design for cost while still achieving the required performance in terms of reliability and availability.

A typical short-circuit current supplied by a generator is represented in **Figure 6**. The short-circuit current is composed of a symmetrical current that is decreasing during the time divided into three periods (sub-cycle transient up to 10 ms, transient up to 250ms, and permanent in a steady state) and, in addition, an aperiodic component Idc decreasing to zero at the end of the transient period. The peak value Ip appears in the first half-cycle (10 ms), and the breaking current Ib is the symmetrical short-circuit value for the shortest circuit break tripping time delay.

Figure 6
Short circuit current wave form.

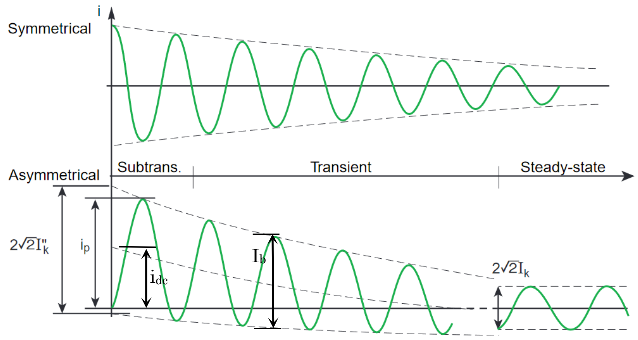

I''_k rms value of the initial symmetrical short-circuit current
i_p short circuit current peak value
$i_{d.c.}$ aperiodic. component of short circuit current
I_b rms value of the symmetrical short-circuit breaking current
I_k rms value of the steady-state short-circuit current

Table 1
Short circuit currents calculation with one engine and 8 engines in parallel

Short-circuit (kA)	I"k (kA rms)	i_p (kÂ)	Calculated at t=55ms				Minimum breaking capacity for IEC Circuit breaker
			Ib (kA rms)	Idc (kA)	Idc%	Ik (kA rms)	
One generator at 15kV	3.65	9.37	2.74	3.22	83%	1.17	3.89
8 generators at 15 kV	29.16	74.73	21.92	25.25	81%	9	30.8

The maximum short-circuit current according to IEC 60909 is calculated for 8 generator units in parallel. The aperiodic component Idc% is calculated according to the formula:

$$Idc\% = \frac{idc}{\sqrt{2}} \times Ib$$

For MV circuit breakers, according to IEC 62271-100, the circuit breaker that opens its pole in 55ms is tested to break an aperiodic component Idc% of 30% maximum. When Idc% is above 30%, a derating is applied on the breaking capacity (see the Schneider Electric MV guide) using the following formula:

$$Minimum\ breaking\ capacity\ according\ to\ IEC = Ib \times \frac{\sqrt{1+2\times(Idc\%)^2}}{\sqrt{1+2\times(30\%)^2}}$$

According to **Table 1**, MV switchboard equipment with the following characteristics are suitable:

- rated peak withstand current of 79 kÂ
- rated short-circuit breaking current of 31.5 kA rms
- rated short time withstand current of 31.5 kA rms during 1 second

Eight generators in parallel is the maximum number you can combine and still be cost effective. More than that and the price increases significantly. The case study is defined by a 70 MW data centre with 7 x 10 MW engines, plus one more unit as a spinning reserve to optimise reliability (when a unit fails, the power plant continues to supply the load with any blackout and black start sequence).

Generator redundancy level

The site is a hyper scale data centre of 70 MW with four buildings, each comprised of four floors (5000 m^2/530,000 ft^2). The generator set comprises seven generators plus an additional redundant unit. To select the correct number of redundant generators, a reliability and availability calculation was performed considering only the generator units without the auxiliary systems and electrical distribution equipment. The calculation considered the failures and also all planned maintenance activities. **Table 2** shows that 7 + 2 redundant units is not enough to reach the availability target of 99.999%. A 7+3 configuration, or an equivalent solution N+2 with an 18 MVA backup grid connection, in order to have a cost competitive solution proved to be the best choice.

Table 2
Availability (%) and Mean Time Between Failure (per year) comparison. Assumptions: only the generator units are considered, the power plant auxiliary systems and electrical distribution equipment are not included)

# of Generators	Availability	MTBF (years)
7+2 generators	99.99173%	10
7+3 generators	99.99992%	557
7+2 generator and 18 MVA grid back-up	99.99999%	3,060

For a more cost-optimised solution, we decided to investigate using 7+2 generators and an 18 MVA grid back-up (only for block1, see **Figure 7**). When a fault occurs in a generator, block 1 switches to the grid and the three remaining blocks will be considered as a set of N+3 redundant.

Electrical topology

The site is divided into four blocks of four floors each (5000m^2/53000ft^2). Each block is designed for a maximum demand of 18 MVA. The electrical architecture is composed of a centralised MV generator power plant with two redundant MV feeders for each data centre block. The power plant redundancy level has 7+2 generators plus an 18 MVA grid backup for Block 1 only.

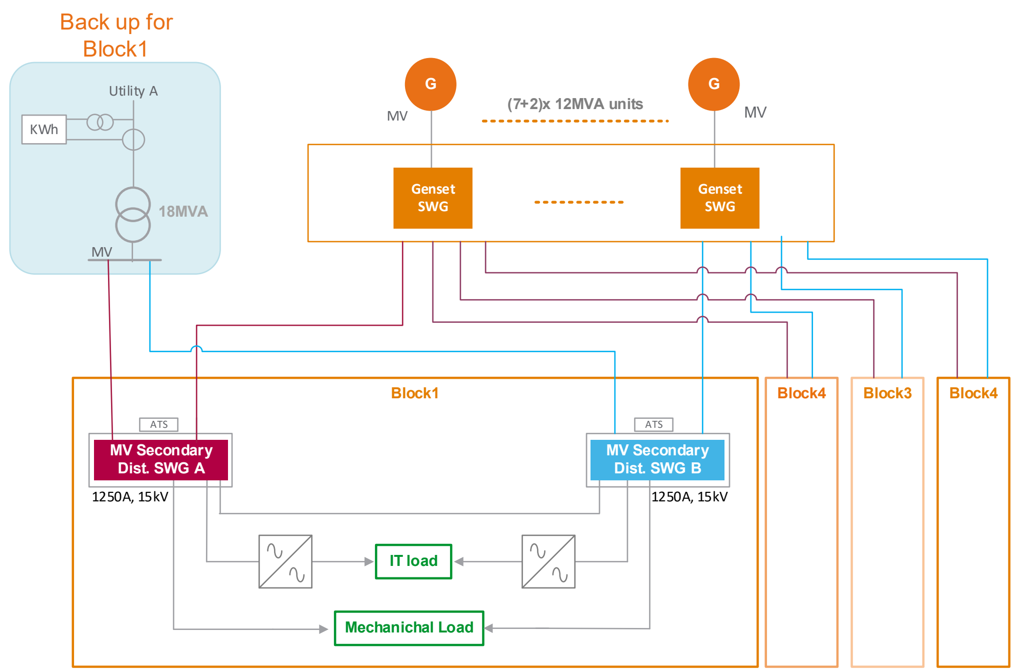

Figure 7
Electrical architecture of the design created during the study to find an optimal solution

This design's easy scalability fits well with the expectations of co-location companies, since they could start with one block to minimise the upfront investment and gradually scale up to the full size over a period of four to five years.

For reliability and fault tolerance reasons, two electrical distribution topologies for the power plant were compared: (1.) a double-fed architecture with automatic reconfiguration, and (2.) a closed ring topology. We decided to choose the latter as it represents a cost advantage in terms of CAPEX (e.g., number of cubicles…) and it has a better protection system based on differential protection with a less complex control system. It is a pragmatic choice considering that less equipment and a simpler control system could enhance the reliability and decrease the probability of having a failure in operation.

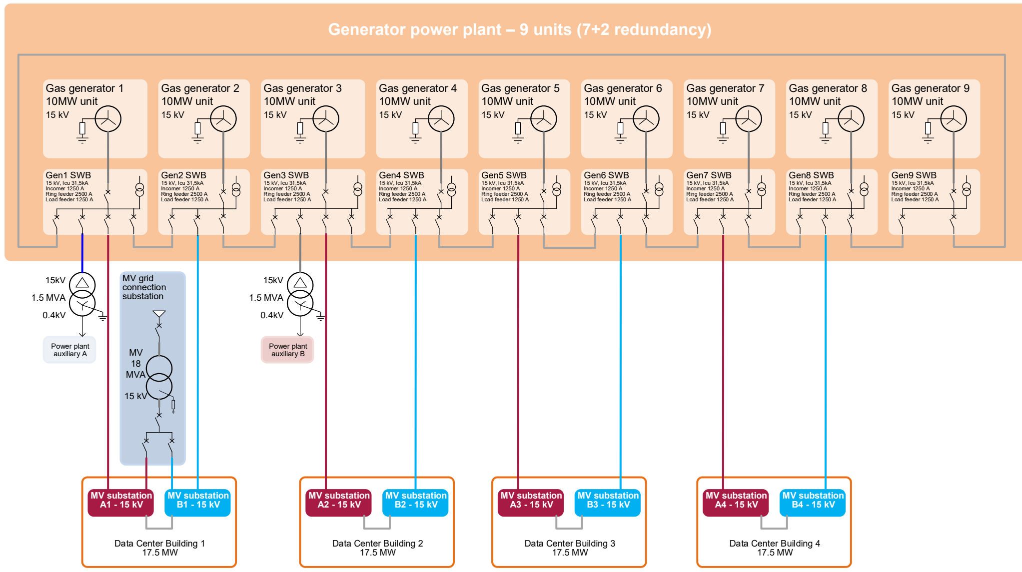

Figure 8
Medium Voltage closed-ring topology

The closed-ring topology with dedicated (by engine) MV switchboards, along with an appropriate protection and control scheme, is a fault tolerant architecture that ensures high availability and, thus, continuity of service for the IT equipment (**Figure 8**).

The protection system is one of the critical pillars of a reliable solution (**Figure 9**). To achieve a fault tolerant system, a protection system based on differential protection (ANSI87) is required. It is based on protecting each zone with differential protection, so that when a fault occurs, only the faulty zone is isolated, and the rest of the power plant keeps running. The bus differential protection (87B), line differential protection (87L), and generator differential protection 87Gen are required.

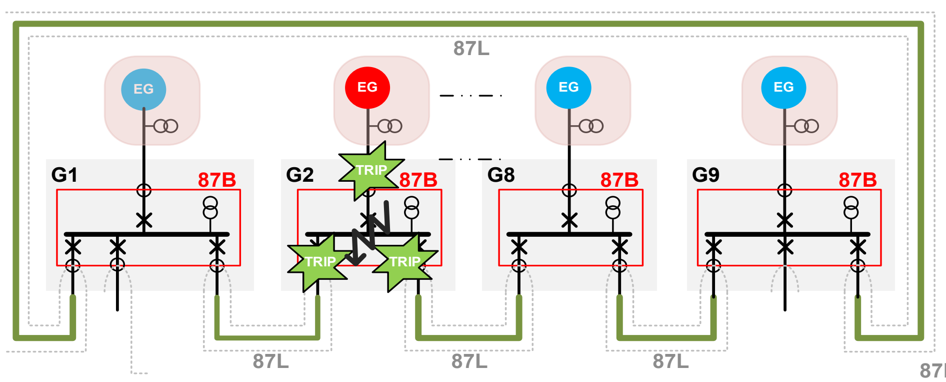

Figure 9
Zoned protection system in a closed loop topology using differential protection, illustration of a fault isolation on the busbar of generator #2

Gas supply system architecture

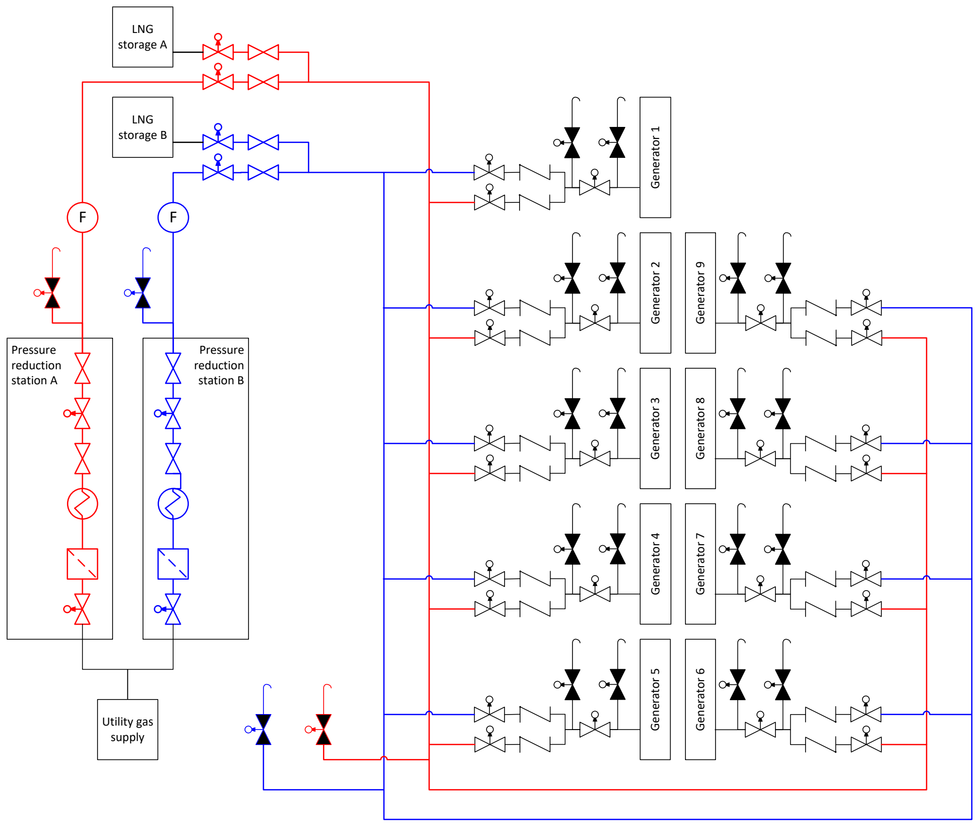

Figure 10
Medium Voltage closed ring topology-based architecture for the gas supply system

The natural gas supply system (**Figure 10**) is designed to be fully redundant to avoid having any single points of failure. It accomplishes this through the following points:

- Utility gas supply is considered the main source of fuel gas
- The site is equipped with a fully redundant (2 x 100%) gas distribution system
- For an unlikely yet possible event whereby the external gas supply is lost, the site has local fuel storage in the form of liquefied natural gas tanks with integrated regasification stations. The tanks would be refilled by tanker trucks or optionally by barges if waterways are available
- Physical separation between the individual engine rooms ensures that any gas leak would require stopping no more than one generating set at a time
- Any section where a gas leak is detected would be automatically isolated.

Generator automation system architecture

Modern gas engine power plant automation systems have evolved from the traditional stand-alone engine speed and load controllers which, although serving the fundamental purpose of control and regulation, often lack features for more sophisticated asset supervision and diagnostics. Indeed, control, protection, and comprehensive supervision features are all incorporated into a modern plant control system, which is a hybrid combination of a centralised plant automation system and distributed and embedded control logics inter-connected together through segmented communication networks (rings). All the distributed power plant units and modules – even the generating sets – are designed to operate separately and independently according to specific sequences based on simple plant operation status signals.

Since the more mundane tasks are fully automated, the operator can focus on more critical, higher-level decisions. For this purpose, the plant automation system collects necessary measurement and status signals, which are then further processed for diagnostics and visualisation purposes. Visualisation takes place in HMI and SCADA displays at the plant control station. The plant control system interface may also be utilised as a remote gateway for external control and asset diagnostics.

The gas engine power plant automation system shown in **Figure 11** is comprised of autonomous and redundant generating set control panels [1] connected together and to the main plant control panels [2] through an Ethernet network ring. In addition to the generating sets, the common main control panels provide interfaces towards common plant auxiliaries, MV and LV power distribution systems [3], as well as the power plant HMI and SCADA server racks [4] through Ethernet network rings. One or both common control panels can optionally have an interface for external monitoring system [5].

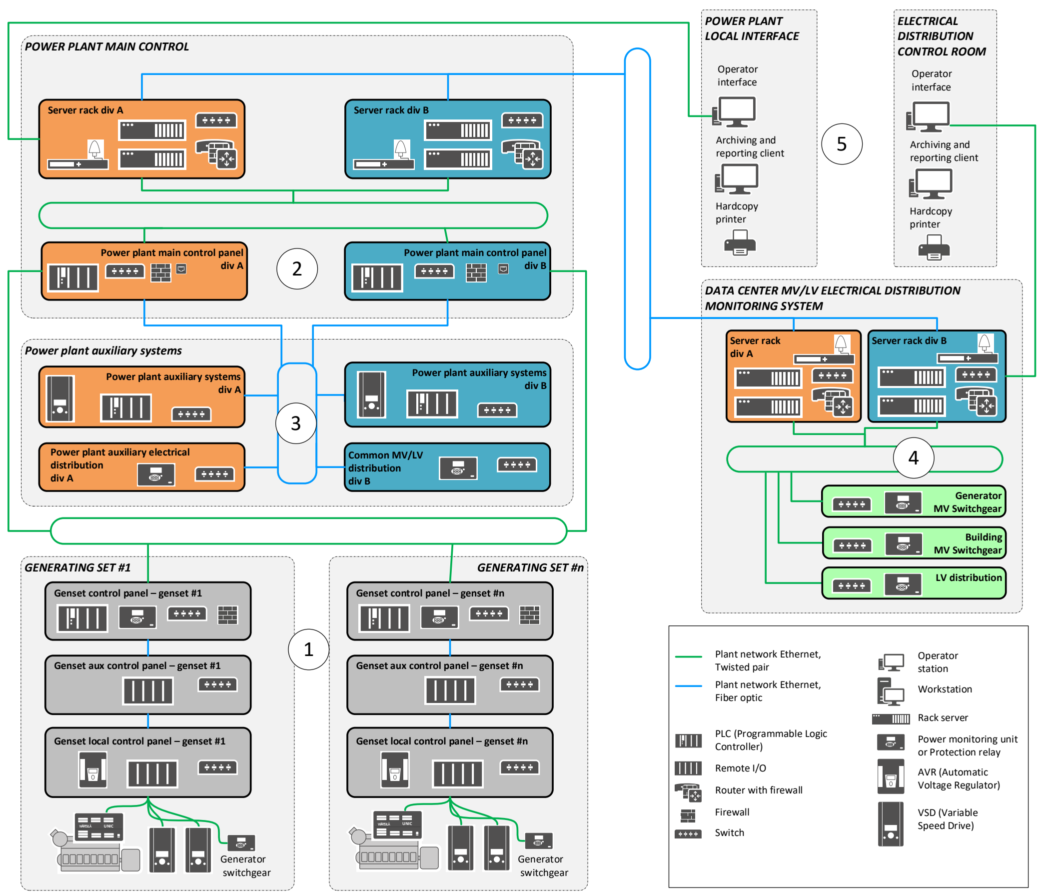

Figure 11
Simplified illustration of the power plant automation system architecture.

Critical auxiliaries power supply architecture

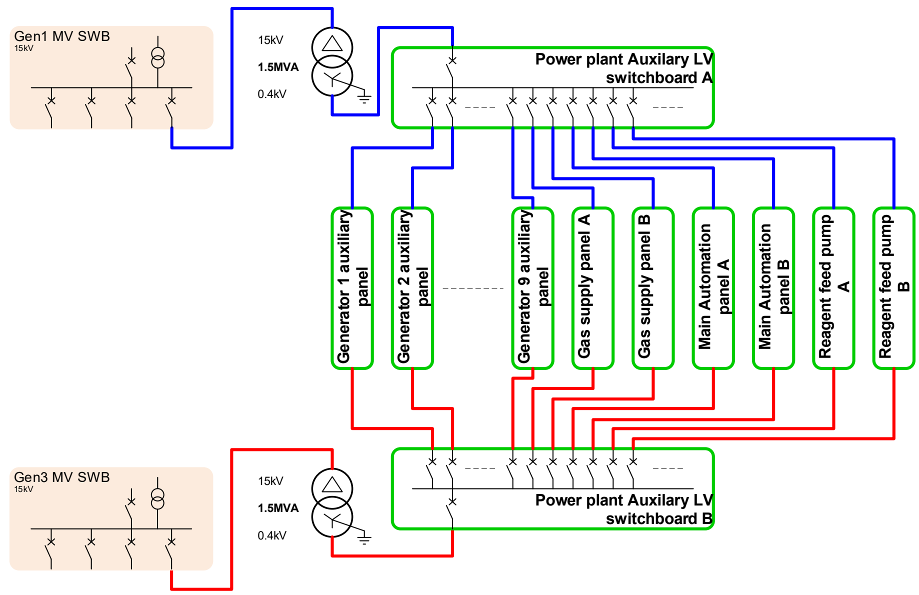

Figure 12
Auxiliaries power supply architecture

The auxiliaries[1] of the critical common systems have fully redundant 2N power supply paths to guarantee continuity of service, as shown in **Figure 12**. Engine auxiliaries, gas supply panels, the automation panel, etc.…are fed from two different power paths: the switchboard of generator 1 being the primary path, and the switchboard of generator 3 as a redundant path to enhance the reliability and ensure continuity of service of these auxiliaries in case of failure.

Comparison of traditional diesel solution vs. off-grid gas power plant

This section takes a standard traditional architecture using an MV redundant grid connection and MV gensets for backup (**Figure 7**) and compares it to the natural gas solution proposed earlier. We chose assumptions based on a reliability study and TCO analysis.

[1] Auxiliaries refers to the control panels for the input and output for the functional blocks such as the gas exhaust unit, piping units, etc.

Case study: 70 MW data centre

The site is divided into four blocks of four floors each. Each block is designed for a maximum demand of 18 MVA. The total active power of the data center is 70MW.

Alternative #1: Traditional solution

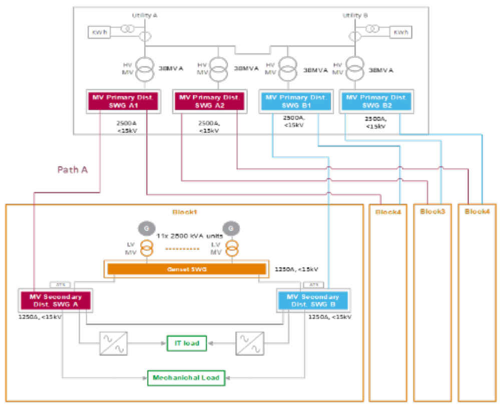

The HV/MV substation is designed to provide redundancy by having:
two redundant HV overhead lines; two redundant sets of 2x38MVA HV/MV transformers; a high voltage Air Insulated Substation (AIS) with a single busbar and a tie circuit breaker; the auxiliaries of the HV/MV substation are designed to provide no common mode failure of the entire HV/MV substation.

The MV distribution for each block consists of:
MV distribution from the HV/MV substation designed with a 2N architecture using AIS technology.
Each block's MV switchboard is equipped with two incomer circuit breakers with an ATS function, nine feeder circuit breakers, and a normally open tie circuit breaker.

The emergency power plant for each block consists of:
11x2MW diesel generators in standby with N+1 or N+2 redundancy (44 engines total for all blocks); monthly load bank tests for each generator.
Single medium voltage switchboard using AIS technology.
Single medium voltage switchboard using AIS technology;
a master PLC managing the main power plant control with hot-standby redundancy;
a single diesel supply system is provided with redundant pumps.

Alternative #2: Continuous gas power plant

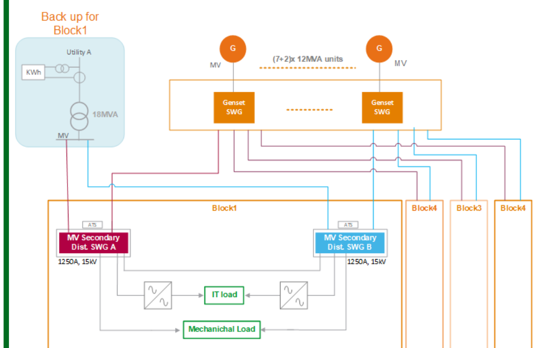

The gas generator power plant consists of:
9 gas generators of 10 MW each; with a maximum of 8 units in continuous operation and 1 unit in standby; All critical common systems (the MV distribution system, auxiliaries power supply, gas supply system, reagent supply system and the automation system) are designed with full redundancy to avoid any single point of failure. The MV distribution for the power plant uses closed ring technology with AIS technology.

Standby MV grid
A standby grid of 18 MVA capacity to improve the generator reliability level

The MV distribution for each block consists of:
MV distribution from the HV/MV substation is designed with a 2N architecture using AIS technology.
Each block's MV switchboard, equipped with two incomer circuit breakers with an ATS function, nine feeder circuit breakers, and a normally open tie circuit.

Reliability comparison

The reliability and availability study presented in this paper aims to show the difference between a traditional architecture with a diesel standby generator power plant and the new architecture with a continuous gas generation power plant, as shown in **Figure 8**. Ensuring high reliability and availability of the server power supplies is the key objective for the data centre's physical infrastructure systems:

- reliability refers to the probability of not having a failure in the supply of power to any IT rack during a given period of time;
- availability refers to the percentage of time the IT racks are powered

Since even the briefest of outages can have severe effects on the business, customers usually define their reliability target as "no risk of power supply failures". Therefore, the customer target is usually more related to reliability rather than to availability. However, some customers hosting their IT process across several data centres are able to handle a single data centre shutdown. They tend to be more interested in the availability of each of their data centres. These customers can formulate their target as "data center availability > 99,999%" (or in the "number of nines"). Considering these customer inputs, the computed reliability and availability indexes are:

- the mean failure[2] frequency (estimated number of failures per year) which reflects the reliability of the data center infrastructure,
- and the mean unavailability (estimated percentage of the time where the data centre is unavailable).

The study is focused on all the equipment from the power sources to the MV switchboards of each block including:

- The grid supply
- The HV and MV electrical switchgear, the MV protection relays and the auxiliary power supply
- The generators including the engine, alternator, control panel, and cooling
- All auxiliary systems, including the fuel supply system, automation, reagent supply system, and the auxiliary power supply

The study estimates the power supply reliability and availability for the server racks in block 2. The "undesired event" shown below in **Figure 13** is defined as "the loss of power to server racks in block 2".

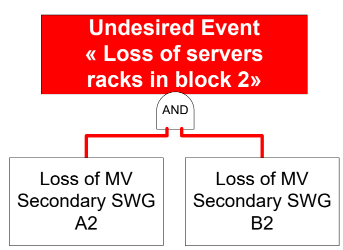

Figure 13
Undesired event illustration

The study is performed using a failure mode and effect analysis (FMEA) combined with a fault tree analysis for multiple contingency analysis. The calculations consider:

- all equipment failure rates and failure modes according to field experience data from manufacturers, as well as from Schneider Electric's reliability data base handbooks.

[2] Failure is defined as a component or system fault that results in the loss of one block.

- all time elements, including detecting the failure, diagnosing the problem, delivering spare parts, and the time taken to repair
- all scheduled maintenance operations and periodic tests

  common mode failures resulting from an overvoltage caused by a severe lightning impact on the data centre building structure

The equipment reliability and maintenance data used for the reliability calculation are based on data from the field experiences of manufacturers and from reliability handbooks, such as IEEE's Gold book and EIREDA.

Table 3
Results of comparing the mean failure frequency, MTBF, mean unavailability and availability of the traditional solutions with the grid as the main supply, with continuous gas power generation.

Summary of reliability studies		Mean failure frequency (/yr)	Estimated MTBF (yr)	Mean un-availability (h/yr)	Mean availability
Traditional solution 1	80MVA HV grid sub-station and "9+1" standby generators	0.0018	563	0.0136	99.99984%
Traditional solution 2	80MVA HV grid sub-station and "9+2" standby generators	0.0002	5348	0.0075	99.99991%
Continuous gas power generation	"7+2" generators and 18MVA standby grid	0.0017	584	0.0743	99.99915%

The results from the reliability calculation in **Table 3** show that:

- the continuous gas power plant achieves the availability target of 99.999%
- The failure frequency representing the reliability index is equal between the continuous gas power plant and the traditional solution with grid connection and standby diesel generators

Total Cost of Ownership (TCO) analysis

Along with proving the technical viability of a new solution, it is also necessary to evaluate its expected economic performance. Obviously, it is not possible to present a thorough discussion of multiple use cases valid for all areas of the world, within the limits of the present study. What follows is, therefore, an analysis of a selected, relatively generic case, involving the same configuration of equipment as described earlier in the paper.

The analysis is based on certain financial assumptions, which reflect the 2019 situation in Western Europe, with a particular focus on the Republic of Ireland. This is because of the observed market trends to locate large data centre projects in that country, combined with considerable difficulties in obtaining a reliable grid power supply within a time frame satisfactory to the investors. For obvious reasons, a generic study may diverge – even considerably – from actual results valid for any specific site. Nevertheless, it does provide certain useful information, showing whether the costs of the proposed solution are comparable to the traditional arrangements with a grid power supply and diesel generating sets.

A TCO analysis was carried out for a data centre plant with a maximum design load of 67 MW (total electrical load for all on-site consumers) operated for 15 years. It was assumed that the facility would be built and commissioned in phases and operated at partial loads. An N+2 configuration would be maintained throughout the project lifetime, along with a limited grid connection, as described previously. Assumed maximum and actual average electrical loads are shown in **Table 4** below.

Table 4
Assumed maximum and actual loads, and the number of installed generating sets

Year	1	2	3	4…15
Max plant load (MW)	17	34	50	67
Average plant load (MW)	15	30	45	60
Generating sets installed	4	6	8	9
Generating sets operating at any given time	2	4	5	7

Selected economic assumptions are shown in **Table 5**.

Table 5
Selected external economic assumptions

Year	Unit	Value
Electricity procurement cost	EUR/MWh	105
CO2 emission fees	EUR/Mg	25
Lube oil cost	EUR/kg	5
Urea solution cost	EUR/dm³	0.45
Pipeline gas cost	EUR/MWh	35
LNG cost	EUR/MWh	40
LFO cost	EUR/MWh	85

Results of the analysis are shown to the left of **Figure 14** below. For the sake of simplicity, annually accrued costs are considered fixed without inflation etc. A relatively high capital expense (CAPEX) in year 0 is attributable to the assumption that the entire MV distribution infrastructure is built before the first generating sets are commissioned. This makes it easier to add more generating sets as the site scales. The relative CAPEX hike in year two is, in turn, caused by splitting the LNG storage facility into two parts, where the first part is installed in year 0 and the second in year 2. The total CAPEX across four years (0 through 3) is around 64.7 million euro. Once the plant is fully commissioned, annual expenditures are at the level of 51 million euros. The largest component of that CAPEX is the cost of natural gas at nearly 41 million euros per year. It is assumed that the plant normally operates on pipeline gas, while the small LNG consumption is attributable to inevitable boil-off in the storage facility. The structure of costs for years 4 through 15 is shown in **Figure 15.**

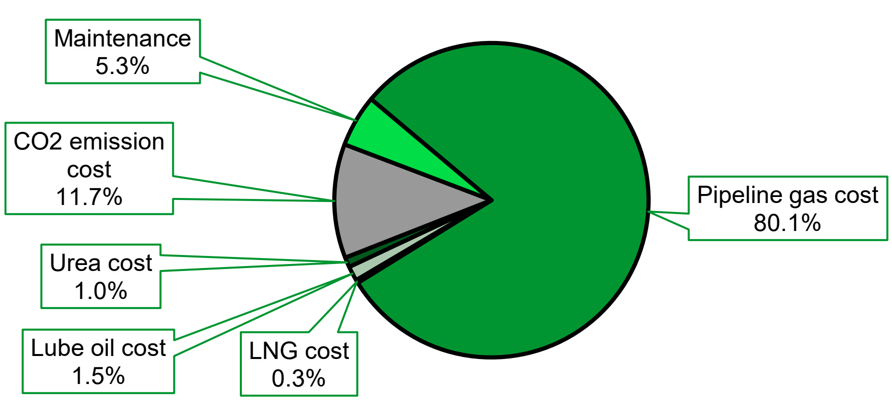

Figure 14
Structure of annual OPEX for the gas plant solution, years 4…15.

The TCO for the on-site gas generation was compared to the TCO of a traditional solution with electricity supplied from the public grid and backup power generation ensured by on-site, high-speed diesel generators. The architecture and assumption of the traditional solution is the same as described in the previous section. For the sake of simplicity, the cost of fuel oil attributable to diesel generating sets was omitted. The TCO results for the grid-diesel solution are shown in **Figure 15**. Naturally, the largest cost component is the purchase of electricity – from year 4 on it is 55.2 million euro per annum. The total CAPEX in this case is 44 million euro.

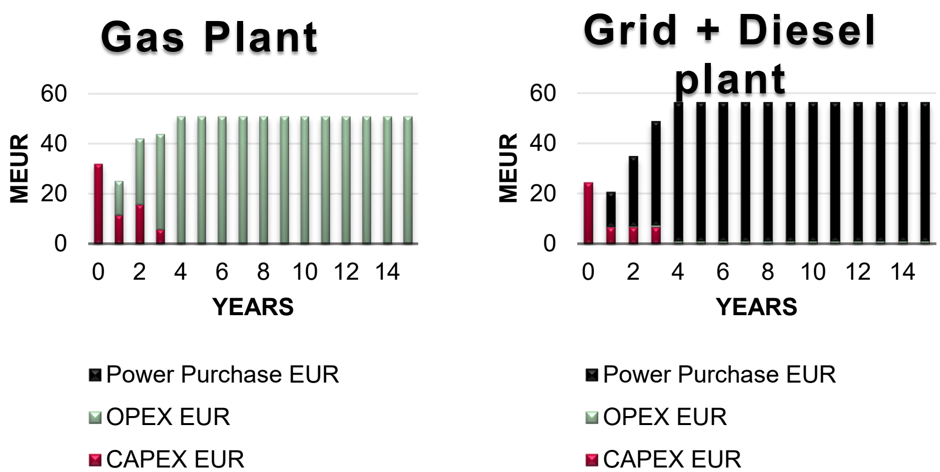

Figure 15
Total cost of ownership for the solution with grid power supply and high-speed diesel backup generating sets.

A comparison of cumulative costs for both solutions is shown in **Figure 16**.

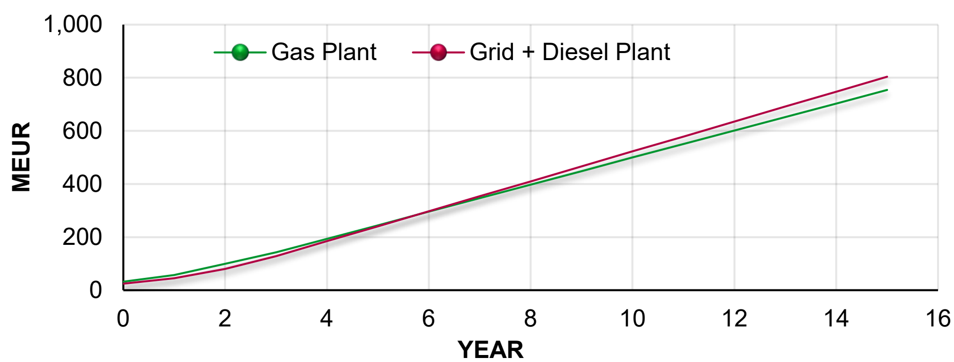

Figure 16
TCO comparison for both solutions

The analysis leads to the following conclusions:

- The TCO for both options is relatively similar. While the gas plant solution is shown to be somewhat less expensive in the long run, the difference does not exceed the accuracy of the assumptions made.
- The gas plant has somewhat higher initial expenditures, attributable to the higher cost of gas engines and related infrastructure, but has considerably lower annual operating costs, because on-site gas generation is less expensive than buying electricity from a public grid.
- It is not possible to clearly indicate which solution is more favourable without accurate knowledge of project-specific conditions.

Note, the analysis was made with assumed conditions relatively unfavourable for the proposed gas generation solution and, therefore, may be considered conservative. It assumes that the gas plant only serves the data centre's IT load. The excess capacity of the engines, both those in operation and the two redundant sets, are not used to generate any revenue (see next section). In practice it might be possible to use those sets to export some power to the power grid, or possibly to supply other co-located facilities (industries, offices etc.), or to provide ancillary services to the power grid. To do this, however, extra expenditures might be required in cases of more complex MV infrastructure architectures in order to perform such revenue-generating functions. Other options include heat recovery (supplying local heat consumers or cooling generation by absorption chillers) or a more dynamic utilisation of capacity (procurement of power from the grid, whenever the costs on the intra-day market are low, and self-generating only when the grid power prices are high). Also, the fact that the grid connection might not be immediately available and might require the temporary rental of some transitional power generation capacity – or delay of a project – was not taken into account. **The general conclusion is that even under relatively unfavourable conditions, on-site power generation may be, at the very least, competitive to the traditional power supply scheme for large data centres.**

Revenue generation

Since a gas-fired, on-site power plant is somewhat more expensive to build than a diesel plant meant only for emergencies, it is advisable to utilise the generators for revenue generation. Potential business models for this are discussed below in no particular order.

Self-generation model

If the data centre has a power generation plant that can be operated continuously, runs on an inexpensive fuel, and has a low emissions footprint compared to the grid, then it should be considered for use as a primary source of power (**Figure 17**). This approach could protect the data centre operator against unforeseen changes in the electricity market, high connection tariffs, as well as higher volatility of electricity prices. At the same time in most power systems and regions, it would considerably reduce the carbon footprint of the facility.

This business model can be complemented with heat recovery functions. When operating, an engine-generator set creates a considerable amount of heat which, instead of being discharged to the environment, can be recovered either for district heating or, with the help of absorption chillers, for cooling the data centre itself. The former is naturally a solution for cold environments and the latter for locations where forced cooling is required most of the year. Both options provide additional revenue without affecting the power generation or concurrent maintainability of the facility.

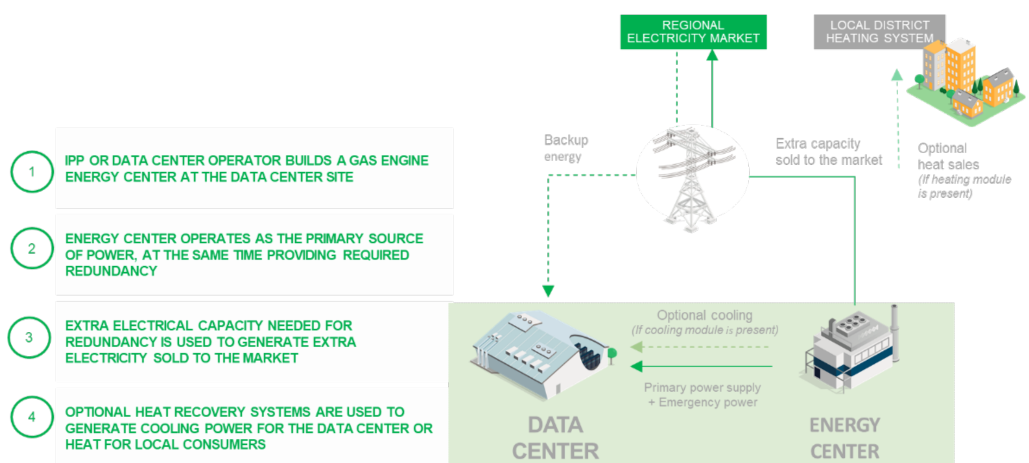

Figure 17
The self-generation model for a gas-fired data centre power plant.

Merchant plant concept

Alternatively, the self-generation plant described above can be used solely as a backup power source for the data centre, while most of the time it would operate independently as a normal merchant generating station, selling its production to external customers. Thus, for most of the time it would act just like a normal gas-engine power plant, while being simply co-located with a data centre. But in the case of disruptions to the grid power supply, it would automatically switch or start-up to supply the data centre. This solution is therefore primarily about getting additional revenue using the new backup power equipment, while keeping the data centre on grid supply.

This alternative might be a preferred approach for markets where the wholesale prices of electricity are very volatile. It enables the profitable operation of peaking power plants, since the engine solution excels with this kind of application. Naturally, the mode of operation can be freely changed at any point in the facility's lifetime, according to changing market conditions. One potential challenge for this type of operation could be the local energy market regulations, which may pose restrictions on using on-site capacity for both self-generation and market generation. This strongly depends on the specific regulations of the energy market in question and needs to be investigated for every project separately.

Deploy power faster than connecting to the utility grid

As larger data centres grow in size and their power demand increases, securing the supply of energy from the grid becomes increasingly problematic. While obtaining a connection with a capacity of a few megawatts is normally not an issue in developed countries, securing a connection of 50 or 100 MW may become challenging, even in urban areas, thus creating a barrier to the development of the data centre. This challenge may be resolved by having an on-site base load power plant to keep the data centre in operation until the grid supply becomes available. After grid connection, the plant might be used in one of the modes described above.

Split ownership

Regardless of the operating model, the plant could be owned by a third party e.g. a utility or an independent power producer (IPP), who would then conclude a relevant long-term agreement with the data centre operator. This solution would allow the data centre operator to focus on the core business without needing to create a division specialised in power plant operations or electricity trading. Outsourcing of the power generation relieves the data centre owner from any front-end investment and turns this over-night cost to monthly fees. At the same time, having a third party owning and operating the power plant would provide a guaranteed flow of revenues, which would help financing arrangements.

Hedging renewables

Gas engine power plants are also a perfect technology for association with local renewable energy sources, such as wind or solar power. Such renewable technologies, although clean and environment friendly, are characterised by their volatile output, and require continuous balancing with dispatchable and flexible energy sources. This makes a flexible gas power plant a perfect companion to allow further penetration of renewable energy sources.

Conclusion

The case study in this paper has shown that in terms of functionality and reliability, on-site power generation for the data centre with medium-speed gas engines is comparable with that of a traditional setup based on a diesel-backed grid supply.

Although the CAPEX of the gas engine power plant is higher, it is compensated with a lower OPEX, which makes the investment more profitable after the first few years. The case study also introduced a concept of a limited grid connection, in which the data centre would be connected to a local MV grid instead of the HV grid. Such a connection would act as a form of reserve supply to compensate for the loss of a single generating set. MV grid connections are typically cheaper and less trouble-some to arrange. Although only the on-site power generation case was studied in this paper, the other numerous benefits of on-site generation with modern medium-speed gas engines can be summarised as follows:

- Natural gas is an inherently cleaner fuel than diesel fuel, and can be cleaner than the fuel mixes for grid power generation in several countries and regions.
- Dynamic performance (i.e. engine startup and loading capability) is almost comparable to the high-speed diesel generators typically utilised for on-site generation, and significantly better than the performance of gas turbines.
- Fuel efficiency is generally higher than that of high-speed diesel engines.
- Low emission levels and high fuel efficiency provide several optional business models to data centre owners, such as a **Self-generation model** for high-cost or volatile energy markets, or a **Merchant plant model** for deregulated and volatile markets. In the latter model the asset owner can earn revenue from grid ancillary services.
- On-site generation also helps in developing and expanding data centre capacity by outpacing construction of traditional transmission and distribution grids.
- Since on-site generation assets can be used profitably, they may be found to be attractive for third parties, e.g. utilities or IPPs. With the help of a third party, the front-end investment would be turned into monthly fees and the data centre owner is relieved of having to operate a power plant. At the same time, the third party gets additional side revenues from the monthly fees unlike in other power plant installations without an adjoining data centre customer.
- Finally, these assets can also be used to support the local grid, which today is facing the increasing penetration of renewable, but intermittent, energy sources, such as wind or solar power.

RATE THIS PAPER ★★★★★

About the authors

Juha Kerttula is a Senior Application Manager of the Data Centre organization at Wärtsilä Energy Business. Juha has studied Electro Mechanics as well as Automation and Control Engineering at Helsinki University of Technology and graduated with a MSc degree in 2001. Before joining Wärtsilä in 2009 he worked in various positions in the energy and mechanical industry, as well as the Finnish Defence Forces. Within Wärtsilä he has dealt with Emergency Diesel Generators for nuclear power plants before entering the Data Centre business segment as a Senior Application Manager. In the Data Centre organization, he is responsible for electrical and automation engineering. Juha is an Uptime Institute's Accredited Tier Designer and is in Vaasa, Finland.

Malik Megdiche received his Master of Science degree and his PhD in Electrical Engineering from the National Polytechnic Institute of Grenoble (INPG) in 2001 and 2004. He is presently a power system engineer and reliability expert at Schneider Electric. His main fields of interest and study are medium and low voltage power system design, protection, simulation and reliability assessment.

Adam Rajewski is an Application Manager for the Data Centre organization at Wärtsilä Energy Business. Adam is a graduate of the Warsaw University of Technology, where he studied Power Engineering. He joined Wärtsilä in 2008, working initially in Sales Support and Sales. Adam has been involved in the development of innovative gas engine power supply solutions for data centres since 2015. In the Data Centre organization, he focuses on mechanical systems and process engineering. Adam is an Uptime Institute Accredited Tier Designer and is located in Warsaw, Poland.

Wedian Youssef, received a bachelor's degree in electrical engineering in 2007 and a PhD in signal processing in 2011 at the University of Lyon in France. In 2012, she was an assistant professor of signal processing. Since 2012, she has worked as a multi-technical engineer at Schneider Electric working on multiple projects, including electric vehicle charging stations, metering algorithms, and electrical design optimisation for large data centre applications.

RATE THIS PAPER ★★★★★

Resources

[Browse all white papers](whitepapers.apc.com)
whitepapers.apc.com

[Browse all TradeOff Tools™](tools.apc.com)
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm

Appendix

Acronyms table

N+1 architecture	**Form of resilience, N components have at least one independent back up component**
2N architecture	Form of resilience, fully redundant , mirrored system
MV	Medium Voltage
LV	Low Voltage
Isc	Short Circuit Current
MTBF	Mean Time Between Failure
Rpm	Revolutions Per Minute
LNG	Liquified Natural Gas
LFO	Liquified Fuel Oil
UPS	Interruptible Power Supply
87B	87B Bus Current Differential Protection
87G	87G Generator Current Differential Protection
87L	87L Line Current Differential Protection
HMI	Human Machine Interface
SCADA	Supervisory Control And Data Acquisition
TCO	Total cost of Ownership
Capex	Capital expenditures
Opex	Operating expenses
IPP	Independent Power Producer

# of Generators	Availability	MTBF(years)
7+2 generators	99.99173%	10
7+3 generators	99.99992%	557
7+2 generator and 18 MVA grid back-up	99.99999%	3,060

Detailed information for the reliability study

The tables below detail the reliability study results.

Traditional solution 1		**80MVA HV grid substation and "9+1" standby generators**
Unavailability (h/yr)	Failure frequency (1/yr)	
0.0136	0.0018	
Contribution to unavailability	Contribution to failure frequency	Main contributors
46.9%	0.5%	Common mode failure due to a severe lightning strike
6.70%	0.79%	Failures on MV path A & B
0.3%	5.2%	Failure(s) on the HV/MV substation & Common mode failure on the generator power plant
45.2%	92.7%	Failure(s) on the HV/MV substation & Failures of several generators

Traditional solution 2		**80MVA HV grid substation and "9+2" standby generators**
Unavailability (h/yr)	Failure frequency (/yr)	
0.0075	0.0002	
Contribution to unavailability	Contribution to failure frequency	Main contributors
85.1%	4.7%	Common mode failure due to a severe lightning strike
12.19%	7.46%	Failures on MV path A & B
1.0%	52.4%	Failure(s) on the HV/MV substation & Common mode failure on the generator power plant
1.4%	34.5%	Failure(s) on the HV/MV substation & Failures of several generators

Continuous gas power generation		**"7+2" generators and 18MVA standby grid**
Unavailability (h/yr)	Failure frequency (/yr)	
0.0743	0.0017	
Contribution to unavailability	Contribution to failure frequency	Main contributors
8.6%	0.5%	Common mode failure due to a severe lightning strike
1.16%	9.71%	Multiple failures of generators
77.0%	21.6%	Multiple failures on the electrical distribution
12.7%	67.3%	Multiple failures on the auxiliaries

# of Generators	Availability	MTBF(years)
7+2 generators	99.99173%	10
7+3 generators	99.99992%	557
7+2 generator and 18 MVA grid back-up	99.99999%	3,060