



THE APPLICATION OF GENERATORS IN A DATA CENTRE

Geoff Halliday, WB Power Services Ltd
September 2022



THE APPLICATION OF GENERATORS IN A DATA CENTRE

Introduction

For those outside of the day to day working of this sector, the size and scale of Data Centres are an unknown quantity. Data Centres are a 24/7 operation, forming a critical part of the global infrastructure and varying greatly in terms of physical size, capacity, ownership and resilience.

Business sectors span global internet brands, venture capital companies, telecom providers and financial institutions, to small scale operators. Users are also mixed and include global web service organisations, co-location hosting providers, financial institutions and streaming services through to lowest cost internet service providers. Many are placed at the hubs of the internet, offering a wide and varied range of ICT services that make up a key part of our everyday lives which we take for granted. With most having backup power systems provided at the M&E infrastructure level, a few others secure their redundancy at the ICY level instead.



Our consumption of data and use of web-based services has been growing rapidly at typically above 50% a year for over 20 years (+4% a month). This is a level of consistent growth not seen in many other market sectors. During this period, the capability and capacity (per watt) of ICT hardware has also seen significant technical advances with a CAGR increase in excess of 40%. When translated into energy consumption, this probably equates to an increase of around 10-15%, with data centres accounting for at least 3% of global CO2 emissions.

On a global basis, the 'best guess' power consumption has been growing at 10% CAGR. With issues such as proximity to end user business, availability of power and connectivity being key location drivers, the larger data centre builds tend to be focused in a key few locations, including Slough, West London and Hemel Hempstead in the UK, and Dublin in Ireland.

Dealing with the application of a generating set to any requirement can be complex. This is particularly true of data centres. This article aims to cover some of the key elements that need to be considered.

Why are Standby Generators required?

The primary reason standby generators are included within a data centre is to provide back-up power in the event of a prolonged power utility failure. This utility failure can be anything from just a few seconds to hours or possibly days. Most operators will set a minimum operational period before refuelling of the installation is required.

The data centre electrical load off, when viewed very simplistically, can be split into two areas:

1. UPS (Feeding the technical IT loads)
2. Cooling load (mech plant)

The UPS package will typically offer a support time of 5 to 10 minutes which is more than sufficient to deal with any short term “brown out” or the 10 to 15 seconds required to bring the standby generating sets online.

The mechanical plant, and in particular the cooling plant, usually has more flexibility and resilience built in by way of cold-water buffer vessels. These are used to hold cooling water temperatures within tolerance for between 30 and 60 seconds, providing sufficient time for the generating sets to synchronise, ready to accept load and for a managed reapplication of load to take place.



Figure 1 – Cooling infrastructure

Starting Points

The starting point of any design should be with the relevant standard and in this case, it is BS ISO 8528 which was updated in 2018. As with any specification, it only sets out the basic **minimum standards** for the equipment (some manufacturers exceed these) and often doesn't move quickly enough in terms of recognising more cutting-edge product developments. For example, this can be across engine performance or market requirements, and have the tendency to look out of date even when newly published. Two examples are detailed by way of an illustration in this article.

The standard does not require a generating set to be capable of achieving any specific level of first step load acceptance, as this capability is based on the Break Mean Effective Pressure (BMEP) that any particular engine is able to deliver. First step load acceptance is, in many cases, a key design performance requirement (e.g. pump starting).

Although having been improved and refined over many years, several of the large engines used on generating sets in the data centre market have been in production for many years. Those that are of a newer design typically offer many advantages such as lower fuel consumption and as a result, lower emissions. Those improvements though can come at a cost, and this can be the first step load acceptance capability of the generating set.

Whilst the Uptime Institute is not a standard by which to work, if a classification and certification is required then consideration must be given to the requirements stated therein. To achieve a Tier III & IV certification requires the standby generator system to be rated for the maximum site load for an unlimited running time with a provision of generating set redundancy. In Europe, the utility or grid supply is accepted to be the “economic alternative” with the generating sets as the standby system.

Power Rating Categories - BS ISO 8528 – 1; 2018 Section 14.3

For commercial reasons, generating set manufacturers use the same engine (with model variations) across a range of power nodes. In addition, a single model of generating set can carry a number of different capacity ratings i.e. 1000kVA ESP, 1000kVA PRP and 700kVA COP; the set rating is dependent on the operational duty / types of loading applied to the generating set, rather than the maximum engine capacity.

Section 14.3 of the standard covers the five different ways in which a generating set can be “rated”. These are COP, PRP, ESP, LTP and DCP. Some of these ratings, for example PRP (Prime) rated sets, have their power ratings set against the ability to deliver into a varying load. They deliver an average load level over a 24-hour period whilst still being able to provide a 10% overload, one hour in twelve.

Section 14.3.3 of the standard describes its operation as

“...being the maximum power which a generating set is capable of delivering continuously while supplying a variable electrical load when operated for an unlimited number of hours per year under the agreed operating conditions with the maintenance intervals and procedures being carried out as prescribed by the manufacturer.”




Figure 2 – Illustration of PRP / Prime Rating (2)

In the case of the ESP rating, there is no overload capability but there is a limitation on the number of running hours it can be operated over a twelve-month period. It is important to note that the way in which these ratings are allocated will vary from manufacturer to manufacturer, based on engine performance, connected rating of alternator performance and maximum operating ambient temperature.

With this extensive range of generating set capacities and rating options, it can be daunting for the design team to fix the direction of this aspect of the design. For most European markets, the grid is a very stable source of power, and it is very unlikely that power will be off for more than a few hours per year (testing included in this number). This is why in theory, an ESP rated set should be an acceptable option, with many preferring the additional flexibility that the PRP option offers.

Depending on the data centre requirements, any of these ratings are valid. However, if the overall design is to be compliant with the principles of the Uptime Institute, then the generating set package should be rated for unlimited running time at 100% load with the utility used as the 'economic alternative'. In Europe, generally the grid will always be used as the operational source of power. This means that Data Centre Power (DCP) is the rating of generator most widely applied to generating sets used in this market sector with its operational duty cycle being defined as:

"Being the maximum power which a generating set is capable of delivering while supplying a variable or continuous electrical load and during unlimited run hours." (1)

The way in which manufacturers apply the DCP rating to their equipment will vary and does warrant investigation. Depending on the engine, the selected power node and ambient operating temperature, some manufacturers use the ESP rating as their DCP rating. Others will use the PRP rating, and some add a further derating on the PRP rating. This selection can impact initial capital cost, engine fuel consumption and consequence emissions, PUE and general operational costs.

The Kohler KD range typically offers a match DCP ESP rating for European applications making it one of the most compact and cost-effective solutions on the market. (2)

When designing a generating set package for any application, it is important to consider likely maximum operating ambient temperatures required as this does play a big part in how the generating sets are designed and rating capacity determined. In a data centre, the generating set solution must always be rated to supply the peak site load, on the hottest day of the year, unless the client gives clear alternative direction.

First Step load Acceptance

We identified earlier that the BS ISO standard does not require the generating set to achieve any specific level of first step load acceptance, as this is determined by the BMEP that the engine is able to deliver. Hence there is a wide variation between manufacturers and the engines used at various power nodes.

60% first step load acceptance for "any rating" of generating set has for a long time been considered an "industry norm" and is often written into many "standard" consultant specifications. However specific consideration is not given to the actual operating requirements of the infrastructure or being qualified with the essential giving performance classification (G1-G4).




Figure 2 – Mechanical Transient Response (1)

The design of many of the larger engines (1500kVA +) used on the larger generating sets in the healthcare, water treatment or hyperscale data centres particularly have been around a very long time. Whilst they have been enhanced and improved over time (more electronic engine management, improvements in the combustion process, high pressure fuel injection, improved materials etc) there are few truly new engines. Those that are new typically offer lower fuel consumption, hence lower emissions, and have a more compact footprint i.e. higher power density.

Those improvements can come at a cost, and this can be the first step load acceptance capability of the generating set. Whilst these advancements are accepted within the standard, the “industry norms” often do not keep pace with, nor understand, the changes or embrace the advantages that they bring. This is often the case when we look at the data centre environment.

Governing Standards / Performance Class ISO 8528-1: 2018 - 8

The standard lays out four performance classes G1-G4 with G4 being the most onerous. In simple terms the performance classes set out the maximum voltage and frequency deviations permitted on the application of a first set load and the time over which the generating set needs to return to a steady state condition and other operational considerations. It is in this area where it is the author’s view that the standard, whilst seeking to provide an all-embracing minimum standard, is failing to keep up with product development.

With the harmonisation of standards across Europe, and to a lesser extent globally, the performance of electrical equipment generally has improved, particularly in relation to the efficient use of energy and reduction of reinjected waveform distorting harmonics prevalent in nonlinear devices. These older devised are the primary loads of any data centre such as UPS and inverter / soft start drives.

Class G2:

This applies to generating set applications where its voltage characteristics are very similar to those for the commercial public utility electrical power system with which it operates. When load changes occur, there can be temporary but acceptable deviations of voltage and frequency. (1)

Class G3 speaks more about the “connected equipment making more severe demands on the stability and level of the frequency, voltage and waveform characteristics” (1) and then site examples such as Telecommunications and thyristor-controlled loads. Both rectifier and thyristor-controlled loads can need special consideration with respect to their effect on generator-voltage waveform.” (1)

Class G4 similarly refers to “Data-processing equipment or computer systems” * as its examples.

By inference, the standards suggest that a data centre requires the use of performance classes G3 or G4 as it uses “telecoms and data processing equipment or computer systems.” Whilst those statements are

This means that the generator(s) never get close to seeing anything approaching a 60% load step totally negating the need for this measure and the associated G3 or G4 requirements that accompany them.

Fuels and Emissions

This is a subject which weighs heavily on the data centre designer and users' minds due to extensive regulation and licencing issues. For some time, there has been a drive to reduce CO2 emissions. Many of the larger engine / generating set manufacturers are certifying their engines to use Hydrotreated Vegetable Oil (HVO) which offers a reduction in CO2 of up to 90% (over the life cycle of the fuel) over conventional BS EN 590 B7 fuel. As with conventional diesel fuel, HVO needs to be looked after during its life cycle with polishing and cleaning systems. There are likely to be some early availability issues, but its adoption is likely to provide a measure of interim solution.



Some designers and operators have looked to gas as an alternative as it is cleaner in emissions and particulates than diesel. However, on-site storage is a space and safety problem, with having a local source of sufficient capacity and reliability another.

Gas, having a lower calorific value than diesel, means the engine swept volume must increase in order to deliver the same power, meaning higher initial capital costs. Starting a gas engine can take much longer and the set is much slower at accepting site loads. Those issues, combined with higher maintenance costs, have ensured gas generation is very rarely adopted.

Many of the larger manufacturers are currently working on a hydrogen fuel solution, running test programmes to assess long term reliability issues. This option looks to offer a promising future but there is some way to go with work still to be done in order to provide a truly "green" source of hydrogen fuel, workable local safe storage and distribution methods.

From a regulator basis, a data centre operator can confront several different licencing, permitting and regulatory requirements, varying from legal jurisdiction, location to location within a given jurisdiction and size / rating of generator installation etc. Much is covered by the Medium Combustion Plant Directive (MCPD) but other requirements will be locally derived and often dependant on other factors. Due to the relative complexity and local variations available, this article does not seek to cover all aspects just to highlight some key considerations. Some of these are:

- Total capacity of installed generating sets
- Thermal input of generating sets (below 3MWth can bring benefits)
- NOx emissions
- If generators are used for other purposes other than standby i.e. STOR of Grid demand response

In relation to NOx emissions specifically, it is important to study the information provided by each individual manufacturer. The results provided (as with all of the other elements which are present in the gas stream) are measured under specific conditions, including engine load and operating temperature, ambient temperature and distance from engine gas flow. All these parameters will vary when at site, hence it is important that each installation is treated and assessed on its own relative merits. The MCPD indicates a requirement to achieve a NOX level better than 190mg/m3 at 100% load and after a max of 20 mins operation.

It is also important to ensure that the units of measurement and remedy are the same. NO_x levels are usually given in mg/Nm³ @ 5% O₂ at 100% PRP. The established method of NO_x reduction is by the introduction of a Selective Catalytic Reduction unit (SCR) into the exhaust gas flow. NO_x reduction/elimination is achieved by injecting ammonia into the gas flow prior to it passing over the catalyst. The use of a precise closed loop ammonia injection system ensures that there is no discharge of unused ammonia. It is important that the exhaust gases (350 – 450°C) are at a high enough temperature to ensure the catalysation process can occur; this is achieved by good design and a generating set working at load. As the NO_x produced in the engine cylinders directly linked to engine pressure and temperature it should be noted that as the exhaust gases expand and cool then the NO_x levels will fall on their own.




Figure 4 – Diagram of a SCR unit

There are of course other elements present in the exhaust gases emitted from a diesel engine. These include HC, PM and soot. The amount discharged depends on rating of set, engine type, fuel consumption, quality of fuel used and engine load. Each of these elements can be reduced by fitting a specific reduction unit. There is however a limit to what components can be added whilst staying within the engine's operating parameters.

Some Additional Considerations

Enhancing generator starting reliability is key to ensuring the performance of any data centre. Fitting dual starting batteries, dual battery chargers and dual starter motors to each set is a quick and cost-effective way of achieving that objective.

Selecting the right generating set control system is key to the overall operation of the system. Firstly, ensure that the selected control system can provide all the necessary communications required for connection to any BMS/ EMS system and site wide protocols. Getting the generator package (multi sets) up to speed and on load as quickly as possible follows. The use of dead bus synchronising can see this achieved in around 10 seconds from start signal regardless of the number of sets to be synchronised.

As discussed earlier in this article, different manufacturers have their own way of arriving at a DCP rating for a given size of set. It is important to understand this in the context of each operational environment. At one end you can have a situation where the DCP rating is the ESP rating whilst the other end it could be a COP rating. There is a +40% difference in power output / capacity between the two extremes. Combine that with a redundant generator requirement of the Uptime Institute it is easy to end up with sets that are light load running or "wet stacking" as it is often referred too. This is damaging for the engine and causes a build-up of unburnt fuel in the exhaust system. It will also hinder the performance of any SCR that might be fitted.

Right sizing the generator package in a scalable way can mitigate this problem but a good design strategy will ensure that a permanent load bank, both resistive and inductive, is connected to achieve a full load test on each machine and raise the load to >25-30% at near to unity power factor regardless of the actual load across all machines.



Regular and thorough maintenance of all plant is key to a reliable M&E infrastructure. In the case of a generating set(s) this should include all control functions and testing the engine at load.

Acknowledgements, Clarification, References and Bibliography.

- (1) Definition from BS ISO 85281:2018
- (2) The engine configuration and much higher power output capacity per cylinder of the Kohler engine when compared to other engines means that the number of cylinders used is less making the engine much shorter. For example, the Perkins powered FG Wilson P2500 (2500kVA ESP / 2250kVA PRP / DCP 2250kVA) set has a V16 configuration. The equivalent Kohler engine, the KD2500 (2500kVA ESP/2250kVA PRP/2500kVA DCP)) offers this power in a 12 cylinder format. The 12 cylinder format is used through to the KD2800. The philosophy of high-power output per cylinder continues through the range right up to the KD4500

Figure 3 – Information provided by KUP a Kohler Power Systems Company.

Author
G R Halliday
Business Consultant
WB Power Services Ltd



Geoff Halliday started his career as an apprentice working for Square D (later part of Schneider) before moving into the critical power sector where he has now worked for over 40 years, splitting that time equally between both the UPS and standby diesel generation sectors.

During this period Geoff has held several roles ranging from Customer Service Engineer, Project Manager, Technical Director, Sales Director through to Managing Director.

The Critical Power market exposes the individual to a wide and diverse range of market sectors ranging Health Care, Life Science, Water Treatment, Banking and Finance, Military, Manufacturing, Process Control through to Data Centres of all sizes. Drawing on his management skills, product knowledge and vast application experience amassed throughout his career Geoff now enjoys sharing his knowledge with others.