Bowman eTurbo Systems for Datacentre applications

Datacentre outlook

Applications of Artificial Intelligence (AI) are expanding across nearly all industries and the datacentre sector is responding with an enormous transformation to keep up with relentlessly increasing compute demands. This is driving the immediate need for more powerful, efficient and flexible datacentre power infrastructure, with reports indicating that global datacentre electricity consumption will increase by 165% by the end of the decade (compared with 2023)[1].

Power bottlenecks present a major impediment to datacentre growth with grid connection scarcity making most of the headlines. To grow datacentre infrastructure at the accelerating rates required, hyperscalers are needing to adopt novel approaches. One such approach, gathering momentum in the market, is to build dedicated modular Internal Combustion Engine (ICE) based independent microgrids using local natural gas (NG) infrastructure[2].

Bowman's eTurbo systems technologies are perfectly placed to compliment these 'self-generating' modular NG ICE based sites, enabling both datacentre back-up and prime power AI requirements to be met, reducing investment and operating costs, while further enhancing the operational flexibility and resilience required to bridge the near-term energy supply gap from the public grid.

Why ICE's are in demand

Short timelines - Modular ICE based generators can be procured, assembled off site, and deployed sequentially and rapidly. This provides the opportunity to commission and bring each power supply 'building block' online in phases together with the datacentre GPU racks and infrastructure as they too are installed and commissioned. It is realistic to start ramping up power availability and datacentre operation in less than 12 months from the beginning of construction.

Reliability – By utilising many smaller generating assets, the risk to the datacentre operation of a single unit having a reliability issue is relatively small versus using larger generating assets or when powering directly from the grid. This means that the site can be operated with more modest back-up power and UPS requirements than would need to be considered to cover the failure of a large gas turbine or a public grid black out.

Efficiency - Rapidly increasing compute demands are driving up the electrical power density of GPU clusters from 10's to 100's of kW per rack, which in turn drives more aggressive cooling requirements, with 50%[3] of a modern datacentre's energy consumption used for cooling. Traditionally the cooling has been directly fed using electricity from the grid driving whatever cooling systems are being employed, but ICE's offer the option to use the exhaust heat (normally wasted to atmosphere) to drive absorption chillers to provide cooling (as well as any local heating requirements). This 'Trigeneration' approach allows ~90% of the fuel energy to be converted by a NG ICE genset for use as electricity, heating and cooling for the datacentre. This approach has been calculated to decrease the total cost of ownership for a 100MW datacentre in the US by 22% compared to a NG CCGT providing power only[4].

NG ICE advantage over diesel ICE

NG ICE generation has several advantages for datacentres over the more traditionally used diesel:

- **Resilience** - Regions with stable NG grids provide long-duration backup or continuous operation with no refuelling logistics and little maintenance to consider.
- **Operating costs** - Depending on location, fuel costs are significantly cheaper, 2 to 5 times lower per thermal energy unit compared to diesel.

salesteam@bowmanpower.co.uk

- **Emissions** - ICE out pollutant emissions, NOx and particulate matter (PM), are orders of magnitude lower, allowing continuous running even in urban settings with strict emissions limits without the need for expensive aftertreatment systems.
- **Sustainability** - Carbon footprint is significantly lower with 25% reduction in CO2 emissions compared to diesel, and with many of todays NG ICE's being able to run on biogas or sold under the 'H2 ready' banner, CO2 neutrality can be achieved by upgrading the ICE hardware once the fuel becomes abundant.

There is a downside however, the start and load acceptance times for typical high efficiency high Brake Mean Effective Pressure (BMEP) NG ICE gensets is measured in minutes and transient load response is sluggish resulting in unsatisfactory voltage regulation. As a result, significantly more UPS support for blackout to online transitions (for datacentre back-up), or during the load swings associated with AI model training cycles is required than has been the case for their diesel equivalents (which can perform the same manoeuvres in seconds and / or with better frequency and voltage stability).

NG ICE technology options

OEM's have been taking steps to improve the start times and block load acceptance capability of their NG ICE offerings to try to keep datacentre UPS requirements modest. Low cost / low development effort options include using generators with low inertia (good for fast start up to rated speed) or high inertia (good for prime power / block load capability), decreasing the ICE boost pressure requirement (increased volumetric efficiency / reduced Miller cycle), reducing the turbocharger inertia (smaller less efficient turbos) and / or increasing governing reserves within the turbocharger match (downsizing turbine nozzle area) while adjusting spark timing and / or decreasing piston compression ratio (CR) to maintain Methane Number (MN) capability. These changes are successfully enabling NG ICE participation in datacentre back-up power applications, but it is evident that the efficiency, power, or both, are compromised for these 'Fast start' versions, versus that achieved with the baseload high efficiency versions that have been developed within the same ICE series.

Additionally, OEMs are deploying other technologies to improve start and loading capability while reducing voltage swings to provide the power quality and stability required during datacentre operation. Options in the market today are to provide faster fuelling response (High Pressure Fuel Injection (HPFI) [5]) and the use of Kinetic Energy Storage Systems[6], [7] and large flywheels to maximise the spinning inertia to bridge the power gaps during power outages or to reduce frequency and voltage dips during aggressive prime power load changes. Other options which could be considered include variable geometry turbines, variable valve timing or electrified turbocharging.

Bowman eTurbo systems assessment – Datacentre back-up

A baseload NG high efficiency high BMEP 60Hz ICE[8], and a 45 seconds 'fast start' 60Hz derivative[5] are compared in the following tables and plotted against applying either Bowman's production release StartIQ™ eCompressor system or a Bowman development eTurbocharger upgrade to the baseload ICE, with and without the adjustments necessary to enable fast fuel response and extreme ambient capability without derate.

The gensets have been duplicated as necessary to supply back up power to a fictional 200MW datacentre site with parameters normalised to the high efficiency baseload version for comparison purposes.

The modifications for the OEM fast start ICE to achieve the dynamic response and high ambient temperature capability allowing back up power participation in global datacentre markets, come at the expense of increased genset installation cost and fuel consumption of approximately 9 and 5% respectively, when compared to the baseload ICE at site level.

salesteam@bowmanpower.co.uk

	OEM baseload version	OEM fast start version
WG used	N	Y
HPFI used	N	Y
2 x Bowman StartIQ™ eCompressor systems used	N	N
2 x Bowman eTurbocharger upgrades used	N	N
Normalised* electrical power (genset)	1.000	0.928
Normalised* specific fuel consumption (genset)	1.000	1.047
Normalised* installed cost (genset)	1.000	1.010
Time from start command to 100% genset load	120	45
Suitability - fast grid balancing	N	Y
Suitability - datacentre backup (cloud / enterprise)	N	Y
Suitability - datacentre prime load (AI training)	N	N
100% load achievable at extreme ambients	N	Y
Number of Gensets required for 200MW site	60	65
Normalised* installed cost (site)	**1.000**	**1.094**
Normalised* fuel costs (site)	**1.000**	**1.047**

*Normalised to High efficiency baseload genset

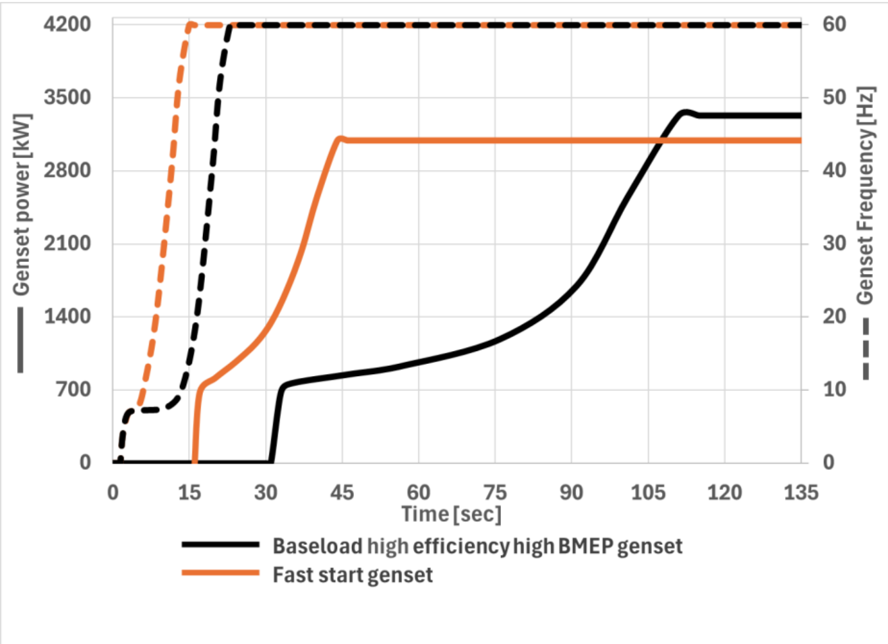

Bowman's StartIQ™ system is an electrically driven high-speed 2-stage radial compressor capable of delivering 0.4kg/s air to the ICE intake or exhaust manifold within 1 second of demand. The system has been successfully applied to multiple high-efficiency, high-BMEP NG ICE gensets without the need for modification, or any change to genset power or efficiency. Load ramps of < 12 seconds from 0 to 100% load with pre-heated ICE conditions (i.e. the ICE has not been run for 24hrs) have been achieved on all NG ICE versions tested to date[9]. Additionally, it has been proven that fast load ramps can be achieved without the need to enrich fuelling / reduce lambda, keeping real-time ICE NOx emissions below that legislated even for steady state running at 100% load, an important consideration for sites with continuous emissions measurement and compliance, or site permitting restrictions. This technology is proven, has been operating for more than 2 years on NG peaking sites supporting the new UK NESO positive quick reserve market[10] and launched into series production in 2024.

The startup and load ramp benefits realised by applying the StartIQ™ eCompressor can be enhanced with an eMachine integrated directly into the ICE's turbocharger (eTurbocharger). A new degree of freedom for ICE control is also unlocked, providing the ability to apply torque directly to the eTurbocharger shaft at any time either for boost assist or by taking torque directly from the shaft to generate additional electrical power

[11].

With Bowman's 6th Generation Power Electronics (PE) capable of 10ms response times to go from 0 to full torque, or vice versa, response times can be realised that are an order of magnitude faster than possible with today's throttle or bypass actuators. This not only unlocks better dynamic performance but also opens the opportunity to control the ICE's boost pressure and, therefore genset load, directly with the eTurbocharger.

By generating electrical power directly from the eTurbocharger shaft, the eTurbocharger speed can be reduced and precisely controlled to deliver the exact amount of air needed for combustion, eliminating the need to use the throttle and / or compressor bypass valves and the losses associated with doing so. By derating the genset load to offset the additional electrical power being generated by the eTurbocharger, fuel savings can be realised compared with the high-efficiency baseload version. Bowman have been developing eTurbocharger solutions together with development partners for industrial markets, with the first eTurbocharger completing its full development cycle and launched into series production in 2025.

	OEM baseload version	OEM fast start version	Baseload version +StartIQ™	Baseload version + Bowman eTurbo
WG used	N	Y	N	N
HPFI used	N	Y	N	N
2 x Bowman StartIQ™ eCompressor systems used	N	N	Y	N
2 x Bowman eTurbocharger upgrades used	N	N	N	Y
Normalised* electrical power (genset)	1.000	0.928	1.000	1.000
Normalised* specific fuel consumption (genset)	1.000	1.047	1.000	0.978
Normalised* installed cost (genset)	1.000	1.010	1.026	1.043
Time from start command to 100% genset load	120	45	37	37
Suitability - fast grid balancing	N	Y	Y	Y
Suitability - datacentre backup (cloud / enterprise)	N	Y	Y	Y
Suitability - datacentre prime load (AI training)	N	N	N	N
100% load achievable at extreme ambients	N	Y	N	N
Number of Gensets required for 200MW site	60	65	60	60
Normalised* installed cost (site)	**1.000**	**1.094**	**1.026**	**1.043**
Normalised* fuel costs (site)	**1.000**	**1.047**	**1.000**	**0.978**

*Normalised to High efficiency baseload genset

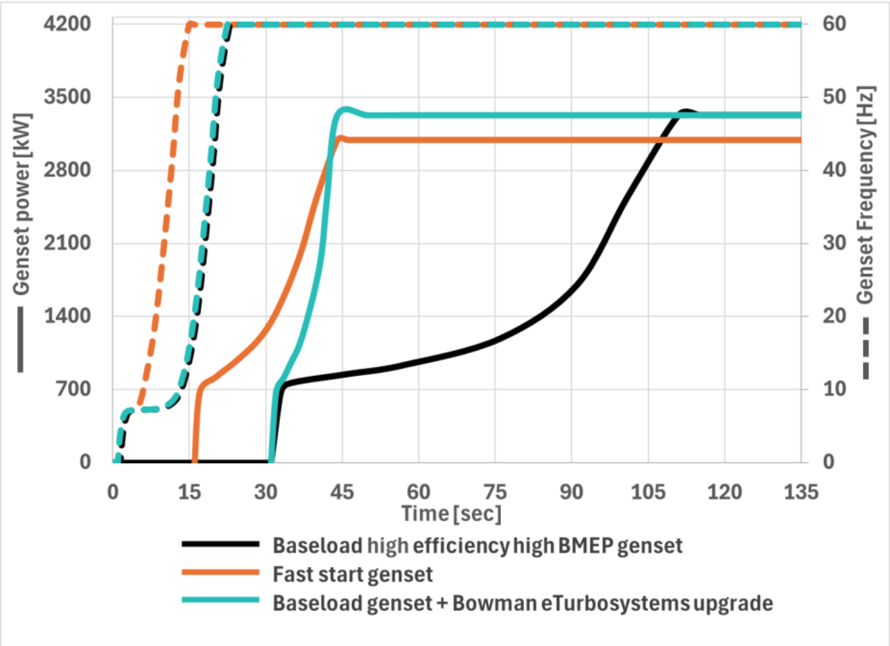

 16/02/2026

salesteam@bowmanpower.co.uk

The addition of Bowman's StartIQ™ eCompressor system to the baseload ICE, with no other changes (i.e. no reduction in Miller cycle or modification to the genset's turbocharger), enables 37 seconds start to full load to be achieved without the need to derate power or fuel consumption. When compared to the baseload ICE at site level, a more modest increase in genset installation cost of ~3% but with 0% increase to fuel consumption is achieved (versus 9% and 5% respectively for the OEM fast start version).

Applying a Bowman eTurbocharger upgrade to the baseload ICE, with no other changes, similarly enables a 37 second start and loading sequence, but with the added ability to generate electrical power directly from the turbocharger shaft such that fuel consumption savings can be unlocked. A reduction in fuel consumption of ~2% can be achieved with a 4% increase in genset installation cost, when compared to the baseload ICE at site level.

Additionally, applying Bowman's StartIQ™ eCompressor system or a Bowman eTurbocharger upgrade to the baseload ICE, together with adjustments assumed necessary to enable fast fuel response and extreme ambient capability without derate have been compared.

In this case a HPFI upgrade and a reduction in the turbocharger turbine nozzle area (together with a piston CR adjustment to maintain MN capability) have been considered, enabling fast fuel response for reduced start command to 100% load times and year-round running without derate even in locations with extreme ambient temperatures, up to 55°C. This combination enables the start to full load time to be further reduced to less than 30 seconds, reducing UPS requirements, while improving dynamic performance and voltage stability for load steps.

The turbocharger nozzle area and piston CR adjustments result in increased pumping losses and a reduction in the effective expansion ratio. The net effect of both is a decrease in electrical efficiency and an increase in the energy going to the exhaust gas versus the original high efficiency baseload version.

In the case of applying the eCompressor, the pumping losses can be partially reduced by applying a wastegate (WG) to reduce ICE exhaust back pressure to the minimum required to be able to govern the ICE using either the throttle or CBP valve.

In the case of applying an eTurbocharger, a WG is not necessary. By generating electrical power (approximately 2 x 100kWe in this scenario) and directly reducing the eTurbocharger speed such to govern the ICE load, actuator governing reserves and losses are eliminated, and most of the pumping and expansion ratio losses associated with the turbine nozzle area and piston CR reductions are offset. Overall net fuel savings can still be realised compared to the baseline high-efficiency baseload version.

	OEM baseload version	OEM fast start version	Baseload version +StartIQ™	Baseload version + Bowman eTurbo	Baseload version +StartIQ™ + fast fuel response + extreme ambient	Baseload version + Bowman eTurbo + fast fuel response + extreme ambient
WG used	N	Y	N	N	Y	N
HPFI used	N	Y	N	N	Y	Y
2 x Bowman StartIQ™ eCompressor systems used	N	N	Y	N	Y	N
2 x Bowman eTurbocharger upgrades used	N	N	N	Y	N	Y
Normalised* electrical power (genset)	1.000	0.928	1.000	1.000	1.000	1.000
Normalised* specific fuel consumption (genset)	1.000	1.047	1.000	0.978	1.034	0.993
Normalised* installed cost (genset)	1.000	1.010	1.026	1.043	1.036	1.053
Time from start command to 100% genset load	120	45	37	37	**29**	**29**
Suitability - fast grid balancing	N	Y	Y	Y	Y	Y
Suitability - datacentre backup (cloud / enterprise)	N	Y	Y	Y	Y	Y
Suitability - datacentre prime load (AI training)	N	N	N	N	N	N
100% load achievable at extreme ambients	N	Y	N	N	Y	Y
Number of Gensets required for 200MW site	60	65	60	60	60	60
Normalised* installed cost (site)	**1.000**	**1.094**	**1.026**	**1.043**	**1.036**	**1.053**
Normalised* fuel costs (site)	**1.000**	**1.047**	**1.000**	**0.978**	**1.034**	**0.993**

*Normalised to High efficiency baseload genset

The addition of Bowman's StartIQ™ eCompressor system to the baseload ICE with fast fuelling and high ambient upgrades, 29 seconds start to full load is achieved without the need to derate power. This incurs an increase in genset installation cost of ~4%, with 3% increase in fuel consumption, when compared to the baseload ICE at site level.

Applying a Bowman eTurbocharger upgrade similarly enables a 29 second start and loading sequence, but with the added ability to generate electrical power directly from the turbocharger shaft, a reduction in fuel consumption of ~1% can be achieved together with a 5% increase in genset installation cost, when compared to the baseload ICE at site level.

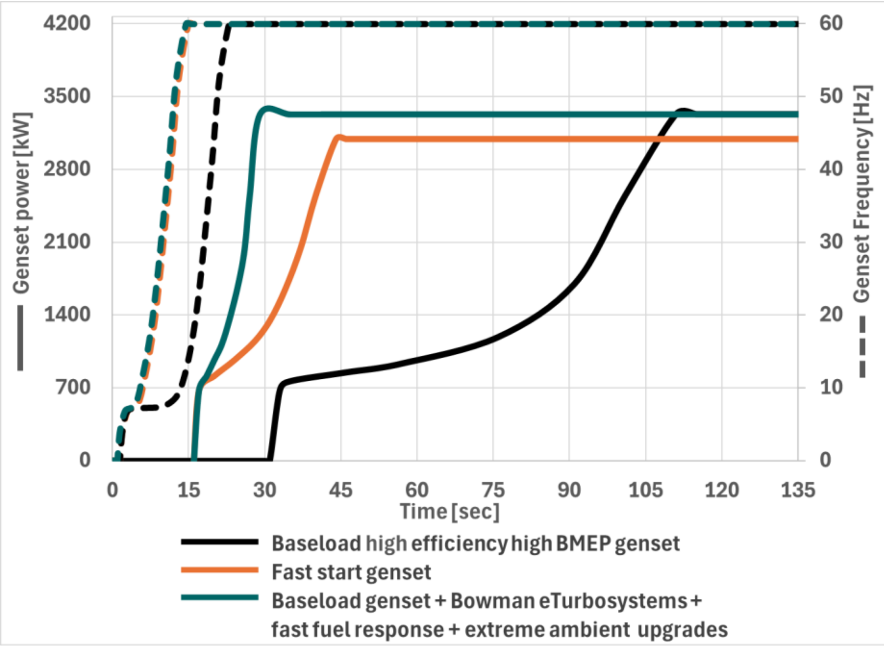

Comparing to the NG ICE fast start version that can fulfill grid balancing and datacentre back up power requirements globally today, the addition of Bowman eTurbo systems to todays best in class baseload versions offers the potential to achieve site level:

- reductions in UPS requirements for black start scenarios of up to 36%
- installed cost reductions of up to 6%
- operating fuel cost (Carbon emissions) reductions of up to 6%

	OEM fast start version	Baseload version +StartIQ™	Baseload version + Bowman eTurbo	Baseload version +StartIQ™ + fast fuel response + extreme ambient	Baseload version + Bowman eTurbo + fast fuel response + extreme ambient
Suitability - fast grid balancing	Y	**Y**	**Y**	**Y**	**Y**
Suitability - datacentre backup (cloud / enterprise)	Y	**Y**	**Y**	**Y**	**Y**
Suitability - datacentre prime load (AI training)	N	N	N	N	N
100% load achievable at extreme ambients	Y	N	N	**Y**	**Y**
Normalised** Start command to 100% load	1.000	1.000	1.000	0.667	0.667
Normalised** Installed cost (site)	1.000	0.937	0.953	0.947	0.962
Normalised** Fuel cost (CO2) (site)	1.000	0.956	0.935	0.989	0.949

**Normalised to Fast Start genset

Bowman eTurbo systems assessment – Datacentre prime power

AI datacentres present unique power management challenges compared to the more traditional cloud enterprise datacentres which are closer to continuous load profiles. During AI training cycles, several thousand GPU's work on the local data in unison, with most of the time spent in power-hungry computation phases (close to the rack's thermal design point (TDP)), and then to a lesser extent at lower power checkpoints or in communication phases (close to idle).

The GPU level power swings, during check points or when switching from computation to communication and vice versa, need to be accommodated, with mitigations possible at either rack level, datacentre level, or both.

Recent methods investigated at rack level[12] include:

1. Software based approaches to inject dummy workloads to fill the power gaps during the communication phase.
2. Firmware restrictions or processing derates to limit GPU power ramp up and down rates and power floors.
3. Rack level energy storage to absorb power peaks and flatten power troughs.

Each rack level method investigated succeeds in that AI datacentre site level performance and reliability can be achieved, but methods 1 and 2 result in excessive energy consumption and method 3 in disproportionate infrastructure costs.

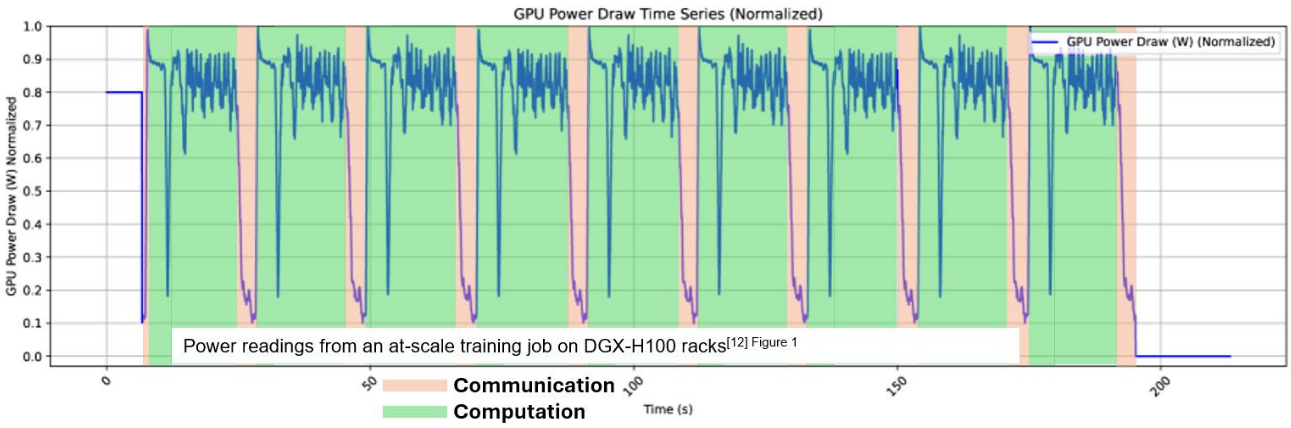

Power readings from an at-scale training job on DGX-H100 racks[12] Figure 1

The peak-to-peak power swings of 80 to 90% of TDP seen at GPU level, translate to peak-to-peak power swings in the region 50 to 60% of the entire datacentre site (reduced versus rack level by ~1/3 reduction as the GPU's are not the only onsite electricity consumer), inflicting significant strain on the local grid infrastructure. Within the computation phase, the high frequency oscillations typically seen in GPU operation are similarly reduced at datacentre level to ~15% peak-to-peak.

The computation phase dominates in terms of GPU and datacentre energy consumption (due to both duration and magnitude of the power requirement versus the communication phase) and therefore any energy savings during this phase, have significant impact on daily fuel costs and CO2 emissions for the datacentre.

Bowman's eTurbocharger solution can help improve the NG ICE performance and transient response through the entire AI training cycle.

- **Computation phase –** The site load is high during this phase and so too is the NG ICE load and thus turbocharger speed. Both, the ICE and Turbocharger are responsive at high loads (i.e. the turbocharger is above the turbocharger lag zone) and both react well to small ICE load changes, ±5 to 10% with minimal frequency or voltage deviations. Provided the inertia of the genset and turbocharger are high enough, the ICE can ride through these high frequency small load variations without the need to take much action.

The addition of the electric motor to the turbocharger has two benefits:

1. A 40 to 50% increase in turbocharger inertia is typical when accommodating a motor between the compressor and turbine wheels, adding permanent magnets and a retention sleeve to the turbocharger shaft. This increased inertia helps to dampen the movements of turbocharger's compressor operating point (improving boost pressure stability and increasing margin to surge) during high frequency site and genset power fluctuations.
2. Significant fuel savings can be realised by generating electrical power with the eTurbocharger during this phase, offsetting most of the losses associated with downsizing the turbocharger nozzle area and adjusting the piston CR to meet high ambient temperatures as previously described. Fuel and CO2 savings of ~4% can be realised. In addition, the 10ms response time of the PE current control (eTurbocharger torque) while in generating mode enables enhanced response to larger frequency or voltage fluctuations should they occur, without the need to react with the throttle or compressor bypass. In fact, the ICE's throttle can be maintained fully open (100%) and compressor bypass fully closed (0%), maintaining the operating point on the compressor map far from the surge line. This maximises the relative movements of throttle and CBP available (to reduce the ICE boost pressure and load with margin to surge) when the next manoeuvre to the communication phase arrives.

- **Transition from computation phase to communication phase –** The turbocharger can be braked during the genset load rejection by generating with the full PE torque available, rapidly decelerating the turbocharger speed. Together with fully opening the compressor bypass valve from 0 to 100% and closing the throttle down from 100% to minimum possible without compressor surge, much more aggressive load rejection rates can be achieved than with a conventional turbocharger together with only actuators available to control it.
- **Communication phase –** During the communication phase the NG ICE genset load would typically be around 30 to 35% to cover the site load with the GPU's idling. Unfortunately, around this load, NG ICEs are at their least capable in terms of accepting load, due to the turbocharger operating at low speed, far from the design point (i.e. both the compressor and turbine operating in low efficiency regions of their respective maps). Accelerating conventional turbochargers to generate ICE boost pressure and load quickly from the communication phase becomes a real challenge i.e. turbocharger lag is at its worst. Two modes of operation with an eTurbocharger are available during the communication phase:

1. Leaving the eTurbocharger in 'idle' mode (i.e. no electric torque is being added to or taken from the shaft), enables the governing reserve to be increased by using the throttle to control the ICE load directly, increasing the ICE's ability to immediately take on load when required (i.e. the pressure and air mass available upstream of the throttle is much more than downstream / in the intake manifold, particularly when matched for extreme ambient temperatures).
2. Motoring the eTurbocharger is also feasible, pushing the turbocharger to a higher operating speed, further increasing and maximising the governing reserves in readiness for the coming load step (throttle closed to its minimum position within compressor surge limits and the compressor bypass opened as required to control the boost pressure and ICE load). This unlocks more possibilities for the upcoming transition to the computation phase, either to take on even more aggressive load shifts or to further reduce voltage and frequency deviations. This eTurbocharger motoring option, comes with a modest energy and fuel consumption trade off versus idling the eTurbocharger, however positive energy, fuel and CO2 savings will be achieved when considered together with the eTurbocharger energy recovery during the longer duration and fuel intensive computation phase.

- **Transition from communication phase to computation phase –** At the beginning of the load shift to computation phase the ICE actuators are immediately saturated (throttle to 100%, CBP to 0%) unleashing boost pressure and air mass flow for the ICE cylinders to consume. Meanwhile the eTurbocharger is motored at the PE's maximum rated current (torque) rapidly accelerating the

salesteam@bowmanpower.co.uk

turbocharger speed to take over the ICE boost pressure and load ramp to meet the site level GPU's demand through to the computation phase.

- **NOx Emissions –** Considering the GPU power draw of the DGX-H100 racks[11], illustrated above, approximately 8 transitions from communication phase to computation phase are undertaken per 150 seconds of the training job. If fueling enrichment is used by the ICE to perform these 8 transitions, significant implications to the average NOx emissions need to be considered. For example, if these 8 ICE transient manoeuvres could take only 3 seconds each, the average NOx emissions over the 150 seconds period would increase by between 1.6 to 2.1 times, depending on how aggressive the enrichment needs to be (i.e. calculation assumes 6 or 10 times the steady state NOx emissions during enrichment respectively).

  It is therefore imperative that the transitions from communication to computation phase (regardless of how long they take) are completed without the use of fuelling enrichment (i.e. zero NOx emissions excursions) when considering applying NG ICE's to prime power applications without expensive exhaust aftertreatment. Note: eliminating fuel enrichment also gives the added benefit of reducing thermal cycling / improving engine durability.

  As mentioned earlier for the datacentre back up scenario, Bowman have previously proven that electrically aiding ICE boost pressure build up with either an eCompressor or an eTurbocharger provides a unique opportunity to perform aggressive NG ICE loading profiles without the need to enrich fuelling or reduce lambda. 7 second, 0 – 100% load ramps have been demonstrated on a high-efficiency, high-BMEP NG genset, with ICE out NOx emissions maintained at ~160ppm through the load ramp and beyond (Note: engine was in hot condition i.e. had been continuously running, and was using modest electrical assist, with eCompressor rating equivalent to ~4.4% of the genset rated power). Without electrical boost assist, the load ramp from 0 to 100% took ~3 times longer, with the peak NOx emissions 22 times higher, and the cumulative NOx emissions through the load ramp period 12 times higher.

  These types of AI training peak load swings have been previously presented by Bowman eTurbo Systems[11], on a similar 22bar BMEP, 1MW NG ICE baseload genset platform. Together with a 20% reduction in turbine nozzle area and a single eTurbocharger applied with a 73kW electrical rating (equivalent to 7.3% of the genset rating), significant positive and negative load transitions could be achieved. The baseline 1MW genset, with enrichment enabled (i.e. NOx emissions > an order of magnitude greater than at steady state) was capable of achieving the load transition from 25 to 100% load in ~ 17 seconds. With the eTurbocharger motoring at full torque (13.5Nm), and no fuelling enrichment, the load transitions were achieved in only ~3.0 to 3.5 seconds (depending on the starting condition of the ICE and turbocharger speed before performing the transition from 25 to 100% load), with NOx emissions complying with steady state legislation throughout.

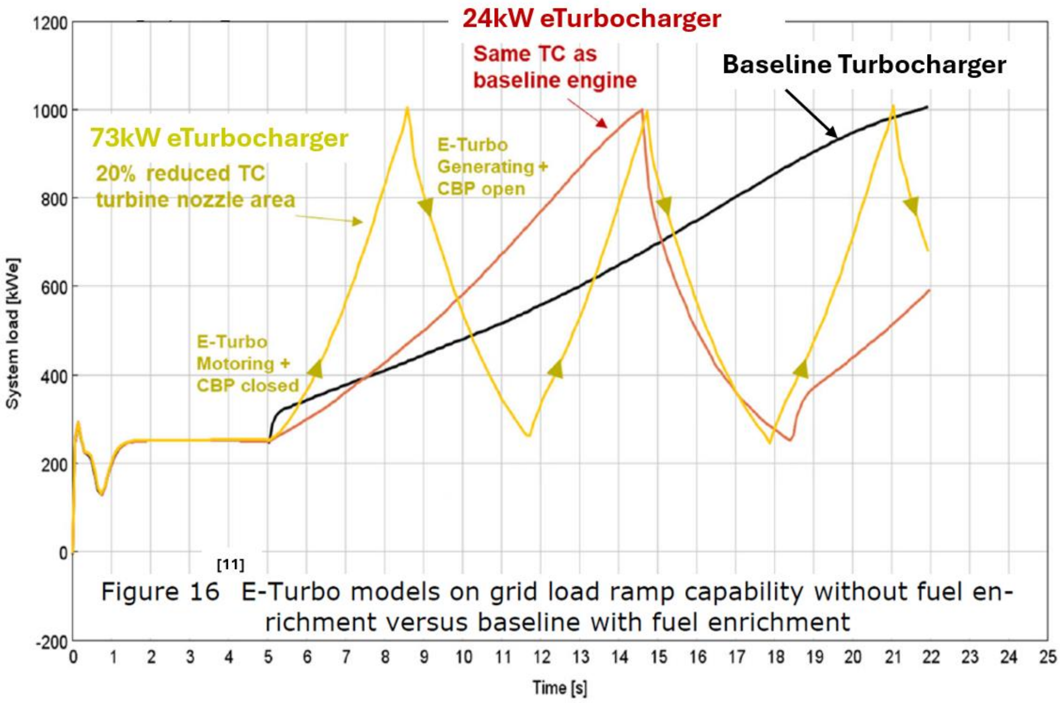

Figure 16 E-Turbo models on grid load ramp capability without fuel enrichment versus baseline with fuel enrichment

For this specific AI training cycle assessment, it has also been assumed that the two eTurbochargers have a combined rating equivalent to ~7% of the electrical power rating of the OEM baseload NG high-efficiency, high-BMEP 60Hz version. Therefore, it is feasible that together with a similar reduction in turbine nozzle area, as required to achieve extreme ambient temperature operation, a 35 to 95% load step could be performed (assumed site level load with GPU's transitioning from communication to computation phase) in less than 3 seconds, with no NOx emissions peaks, ensuring steady state NOx emissions compliance can be maintained through the entire AI training cycle.

This eTurbocharging approach allows the possibility to reliably achieve AI training cycles with steady NOx emissions, without the need for the rack level power smoothing strategies 1 and 3 mentioned above. There will, however, still be a need to modestly use strategy 2. The GPU load ramps will need to be matched to capability of the NG ICE load ramps (~3 second duration) to keep voltage and frequency deviations steady when moving from communication to computation phase and then back again.

To unleash full GPU performance, without firmware restrictions to GPU power ramps, the eTurbocharger electrical rating could be further increased, however, it is not expected that further fuel consumption savings could also be realised by increasing the eTurbocharger rating beyond 7% of the genset rating. An alternative method could be to add further inertia to the genset, which would help minimise voltage and frequency deviations during the seconds while the eTurbocharger deals with the ICE boost pressure and load transitions while maintaining NOx emissions. Therefore, a trade-off needs to be considered at datacentre level, between the cost of sizing the eTurbocharger rating and any increased genset inertia versus fuel consumption benefits and the impact of using firmware restrictions to GPU power ramp up and down rates and on datacentre performance.

Initial calculations indicate a higher inertia 4500KVA generator (inertia of ~ 2.3 x that calculated for the OEM fast start version generator) is sufficient to dampen load variations seen within the computation phase together with fine adjustment of the eTurbocharger torque. Electrically removing or adding maximum torque to the turbocharger shaft to fulfil the ICE boost pressure and load transitions from computation to communication phase and vice versa, the frequency and voltage deviations can be managed during AI load cycles similarly, whether a fast-fuelling response (HPFI) solution is used or not.

The increased genset inertia does add time to the start-up phase of the genset compared to the lower inertia version, but excellent start command to 100% load times can still be realised, with 46 seconds calculated without HPFI. HPFI only becomes a requirement if fast starting is in consideration to reduce UPS requirements for back up scenarios, with 38 seconds start command to 100% load possible. Alternatively, more inertia could be added to the genset (equivalent to increasing the generator inertia by as much as 3.4 times versus that calculated for the OEM fast start version) while still meeting a start command to 100% load ramp time of 45 seconds.

	OEM baseload version	OEM fast start version	Baseload version +Bowman eTurbo + extreme ambient + high inertia generator	Baseload version + Bowman eTurbo + fast fuel response + extreme ambient + high inertia generator
Ratio of generator to ICE inertia	2.6	2.6	6.1	6.1
WG used	N	Y	N	N
HPFI used	N	Y	N	Y
2 x Bowman eTurbocharger upgrades used	N	N	Y	Y
Normalised* electrical power (genset)	1.000	0.928	1.000	1.000
Normalised* specific fuel consumption (genset)	1.000	1.047	0.993	0.993
Normalised* installed cost (genset)	1.000	1.010	1.043	1.053
Time from start command to 100% genset load	120	45	46	38
Suitability - fast grid balancing	N	Y	Y	Y
Suitability - datacentre backup (cloud / enterprise)	N	Y	Y	Y
Suitability - datacentre prime load (AI training)	N	N	Y	Y
100% load achievable at extreme ambients	N	Y	Y	Y
Number of Gensets required for 200MW site	60	65	60	60
Normalised* installed cost (site)	**1.000**	**1.094**	**1.043**	**1.053**
Normalised* fuel costs (site)	**1.000**	**1.047**	**0.993**	**0.993**

*Normalised to High efficiency baseload genset

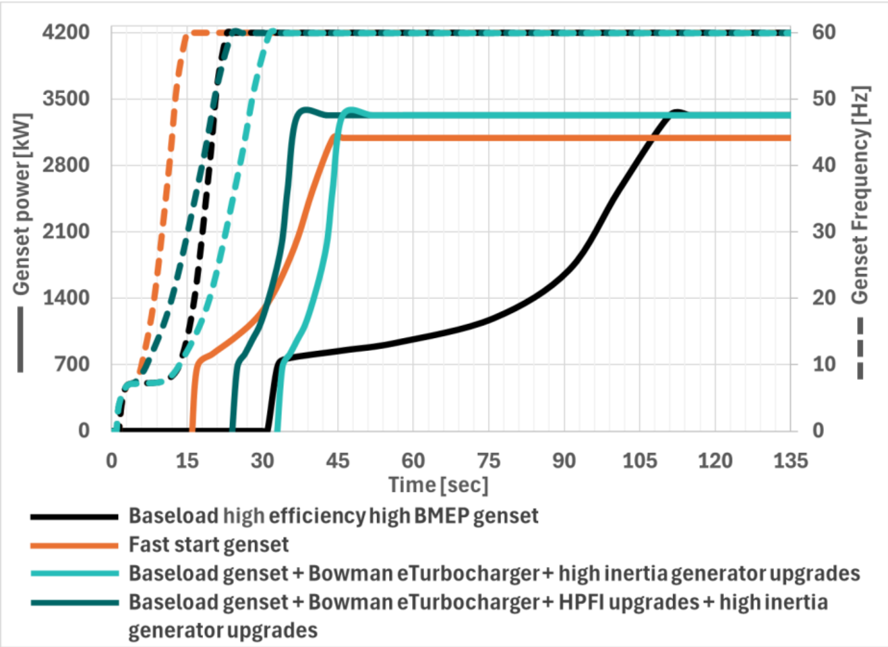

Bowman eTurbo systems datacentre outlook

Considering today's grid capacity and transmission challenges, build out of self-generating modular genset microgrids presents a near term opportunity for datacentres to bridge the energy supply gap. Deployment of Bowman's eTurbo systems together with NG ICE gensets has the potential to address both datacentre backup and prime power requirements, while delivering best in class power densities, efficiencies and continuous emissions compliance, only possible with today's baseload versions running at steady state. This enables enhanced flexibility and resilience of the assets, significantly reduced investment and operating costs, and improved sustainability.

Should the public grid bottle necks and power requirements be addressed longer term, these microgrid solutions can allow for continuous optimisation of datacentre energy costs and economic performance. Continued coverage of multiple scenarios is possible with the same assets without modification, in-house back-up power requirements can be maintained, prime power can still be fulfilled during periods when market prices are favourable, and the business model can be expanded to feed electricity into the public grid by providing reserve capacity for grid-related ancillary services or capacity markets during periods of high pricing. And of course, if the economics are favourable, they could be relocated and re-used in other datacentre locations with similar datacentre power needs and requirements.

References

[1] https://www.goldmansachs.com/insights/articles/ai-to-drive-165-increase-in-data-center-power-demand-by-2030

[2] https://www.constructionowners.com/news/oracle-and-openai-to-power-new-texas-data-center-off-the-grid

[3] https://www.datacentrefrontier.com/sponsored/article/55307885/nautilus-why-liquid-cooling-is-no-longer-optional-for-ai-driven-data-centers

[4] https://www.mtu-solutions.com/eu/en/technical-articles/2025/enabling-data-center-operators-to-bridge-the-grid-capacity-gap-with-self-generation-solutions.html

[5] https://www.jenbacher.com/en/gas-engines/type-6/j620-fast-start/

[6] https://www.piller.com/product/shieldx-dynamic-power-stabilisation/

[7] https://www.mtu-solutions.com/eu/en/applications/power-generation/power-generation-products/mtu-kinetic-powerpack.html

[8] https://www.jenbacher.com/en/gas-engines/type-6/j620/

[9] https://bowmanpower.com/wp-content/uploads/2026/02/CIMAC_Congress2025_Paper-ID-003-Development-of-eCompressor-system-to-improve-high-speed-engine-transient-load-acceptance.pdf

[10] https://www.neso.energy/industry-information/balancing-services/reserve-services/quick-reserve#Technical-requirements

[11] https://bowmanpower.com/wp-content/uploads/2026/02/06_Douglas_Bowman_25.ATK2020.pdf

[12] https://arxiv.org/html/2508.14318v1

salesteam@bowmanpower.co.uk