

THE AMERICAN OIL & GAS REPORTER[®]

SEPTEMBER 2026

The "Better Business" Publication Serving the Exploration | Drilling | Completions | Production Industry



Power Generation Requires Fit-For-Purpose Planning & Installations



Project Specifics Determine Best Way To Generate Electricity

By Parker Gregg

As AI data center demand accelerates and oil field electrification expands through artificial lift, frac operations, electric pneumatics, rigs, and other field applications, the industry is seeing two major power-demand curves converge on the same grid. Data center developers and oil producers are increasingly competing for many of the same resources, from construction labor to interconnection capacity, with little relief in sight.

Goldman Sachs estimates that U.S. data center power demand totaled 31 gigawatts in 2025 and projects that it will climb to 41 gigawatts this year and 66 gigawatts in 2027. Texas Governor Greg Abbott's temporary moratorium on new data center grid connections may soften those projections, but the pause does not change the underlying trajectory: data center electricity demand is on track to double in just two years, a figure that does not even account for the parallel push to electrify the oil field.

Whether power is delivered through grid connection or generated onsite behind-the-meter, access to clean, uninterrupted power supply has become the defining constraint for both markets. Despite

advances in remote power generation, no single technology answers every load profile, site condition, or project timeline. The right answer increasingly depends on speed to power: How fast does a producer or data center developer need electrons?

Oil field power generation has evolved over time. For decades, the industry depended on isolated, fuel-intensive diesel engines driving mechanical equipment, such as pumps and compressors, directly. Diesel gensets eventually took over that role for electric-driven equipment, but stricter air-emissions permitting has since confined diesel to backup and temporary service in most jurisdictions. Fuel costs accelerated diesel's move to a standby role, with natural gas offering a cleaner and less expensive alternative.

Today, natural-gas-fueled generation (small industrial gas turbines and high-speed reciprocating engines) is predominantly powering the electrification of oil field production and facility loads. Data center power is drawing on a far broader mix, including those above as well as aeroderivative, frame-class gas turbines, other reciprocating engines, and not-yet-built combined-cycle capacity. Lesser known, emerging technologies will enter the mix as developers piece together whatever can be delivered on their timelines.

Both markets are now asking a version of the same question: What generation technology fits a given site, timeline, and load profile? Having built and operated a fleet of our own power plants, we do not think there is a single right answer, only trade-offs worth stating plainly. Following, we examine four established paths—simple-cycle gas turbine (SCGT), combined-cycle gas turbine (CCGT), aeroderivative gas turbine, and reciprocating internal combustion engine (RICE)—on their own terms, including where each carries real trade-offs against the others.

Simple-Cycle Gas Turbines

As the name suggests, a SCGT converts fuel into mechanical power or electricity through a single, direct process. Air is compressed, mixed with fuel, ignited, and expanded through a turbine section, with no secondary system to capture or reuse exhaust heat.

Throughout this piece, “SCGT” refers specifically to industrial/frame-class simple-cycle turbines. Technically, aeroderivative units are also simple-cycle machines. The distinction that matters operationally is design lineage (industrial vs. jet-engine-derived), not thermodynamic cycle, which is why we treat them separately below.

SCGT’s single-pass simplicity is the technology’s defining trade-off. Without a water loop to recover exhaust heat, industrial SCGTs typically convert 32% to 38% of fuel energy into usable power. Over time, that lost heat translates into a higher fuel cost per megawatt-hour delivered. What operators buy in exchange is deployment speed, mechanical simplicity, and a machine built to tolerate conditions that would be harder on more efficiency-optimized equipment.

Small industrial SCGTs are engineered for continuous duty first. Heavier construction and wider design margins than aeroderivative equivalents let them absorb variable fuel composition, including associated and field gas, without the tighter tolerances that make some aeroderivative units more sensitive to fuel quality.

This piece focuses on generator-drive applications, turbines producing electric power rather than driving compressors or pumps directly. Mechanical drive, not power generation, is where small industrial turbines have historically logged the majority of their oil field runtime, particularly in pipeline and gathering compression.

That is changing. As field electrification expands to cover artificial lift, drilling loads, and facility power, generator-drive deployment of small industrial turbines has grown into a primary use case in its own right, not a secondary one riding on the compression business. The same ruggedness and fuel tolerance that made these units durable compressor drivers for decades are now proving just as valuable on the power generation side.

On-site maintenance and long overhaul intervals are the norm rather than the exception for SCGTs, and the units are as commonly configured as permanent, prime-power installations as they are for shorter-duration deployment. This is not exclusively rental or bridge equipment; operators evaluating a multi-year compression station or facility load should weigh industrial SCGTs as a first-class permanent option, not a placeholder. Without secondary steam systems, these units can reach full power in roughly five to 10 minutes from a cold start, useful wherever quick response matters even in a fundamentally continuous-duty application.

Where SCGT gives ground is efficiency at partial load and full load alike. Reciprocating engines in the same power class routinely post higher simple-cycle thermal efficiency and hold that efficiency far more consistently as load varies. A turbine’s output and heat rate degrade more steeply below roughly 50% to 70% load than a reciprocating engine’s does.

Ambient temperature compounds a derate that starts with elevation: at roughly 2,700 feet in the Permian Basin, a gas turbine is already operating below its rated output before temperature ever enters the picture, and a 115 degrees Fahrenheit day on top of that elevation can push total output down 40%, a derate many high-speed reciprocating engines largely avoid.

For operators prioritizing steady, continuous, near-full-load operation with fuel flexibility and mechanical simplicity, SCGT remains a strong, proven fit. For operators facing wide load swings, high ambient temperatures, or maximum fuel efficiency as the deciding factor, these trade-offs are worth weighing directly against reciprocating alternatives.

Combined-Cycle Gas Turbines

Combined-cycle pairs a gas turbine with a heat-recovery steam generator (HRSG) that captures exhaust heat a simple-cycle unit would lose and uses it to make steam for a secondary turbine. That second conversion step is what pushes efficiency well above simple-cycle levels. Modern frame-class combined-cycle plants commonly reach efficiencies in the 50% to 60% range, with the most advanced systems exceeding that.

CCGT is, almost without exception, a frame-class technology. Unlike the small industrial turbines this piece otherwise focuses on for SCGT, the scale needed to justify an HRSG and steam bottoming cycle economically pushes CCGT toward the largest turbine frames on the market. That scale is exactly what makes CCGT a poor fit for most oil field applications and a strong fit for large, continuous data center or utility loads, where efficiency compounds into meaningful savings over years of near-constant operation.

The efficiency advantage comes at a real cost in capital, schedule, and flexibility. The total timeline compounds two separate constraints. OEM turbine order books determine how long a producer waits just to secure a slot, a queue that has stretched considerably as demand from data centers and utilities has surged. Once a slot is secured, EPC design, procurement, and construction adds several more years on top of that wait. Combined, total time-to-power for a new CCGT plant can stretch well past five years, beyond what most oil field or fast-track data center projects can tolerate.

For a data center developer racing to bring capacity online in months, CCGT’s lead time alone can rule it out regardless of its efficiency advantage. The result is a technology best suited to operators who can commit capital and schedule years in advance in exchange for the lowest fuel cost per megawatt-hour of any

“Industrial simple-cycle gas turbines typically convert 32% to 38% of fuel energy into usable power. Over time, that lost heat translates into a higher fuel cost per megawatt-hour delivered. What operators buy in exchange is deployment speed, mechanical simplicity, and a machine built to tolerate conditions that would be harder on more efficiency-optimized equipment.”

gas generation technology in this comparison. Where SCGT, aeroderivative turbines, and RICE compete primarily on speed, footprint, and operational flexibility, CCGT competes on decades of efficient, low-cost baseload output once it is finally in service.

Aeroderivative Gas Turbines

Aeroderivative machines trace their design lineage to aircraft jet engines, adapted for ground-based power and mechanical drive. Their heritage shows up directly in performance: lighter weight, higher power density, and simple-cycle efficiency generally better than industrial gas turbines. Aeroderivatives also start fast and reach full load within minutes, including black-start capability in some configurations, offering a real advantage for applications where response time matters as much as steady-state output.

Mobile, trailer-mounted aeroderivative packages built specifically for fast-track and rental deployment have become genuinely common in the oil field, used for temporary compression, bridge power, and emergency response where a project needs megawatts on site in days rather than months. These units' mobility has led many in the industry to assume aeroderivative equals rental, industrial equals permanent. That assumption misses a lot. Larger aeroderivative units are widely deployed as permanent, prime-power installations in their own right, for baseload generation, peaking, cogeneration, and mechanical drive alike.

What carries over from the mobile packages, even into permanent, non-trailer installations, is a comparatively modular approach to installation. Aeroderivative units arrive largely pre-packaged and factory-assembled, requiring less extensive balance-of-plant work than a small industrial or frame-class turbine typically demands. A gas turbine installation still carries real complexity: auxiliary systems, inlet and exhaust ductwork, and commissioning work a reciprocating engine installation doesn't require. Aeroderivative sits between the two, simpler to site than industrial or frame turbines, but still a turbine installation, not a modular engine skid.

The real distinction is not the turbine class but the specific package: a handful of purpose-built mobile products account for most of the rental fleet, while permanent installations across turbine types make up a substantial share of the broader install base.

What operators give up for that speed and density is the ruggedness that defines industrial SCGT. Aeroderivative units generally hold tighter tolerances inherited from their aviation origins, making them more sensitive to fuel quality and composition, and their major overhaul intervals tend to run shorter than an industrial turbine's, a direct consequence of the same lightweight, high-performance design that makes them fast and efficient in the first place.

That sensitivity shows up mainly upstream, in the fuel gas conditioning required to protect the unit. Aeroderivative installations often need tighter gas treatment systems than an industrial turbine's wider tolerances allow, adding capital and complexity to the balance of plant.

For operators who need capacity fast, in a compact footprint, or with the flexibility to relocate equipment as a project's needs change, aeroderivative turbines remain a proven, widely used answer for both permanent and temporary applications.

Reciprocating Engines

Much like an automotive or diesel engine, reciprocating internal combustion engines burn fuel in a series of cylinders to turn a crankshaft connected to a generator. While they are a familiar technology scaled up rather than a purpose-built power-generation machine, they have many strengths.

Efficiency is the clearest advantage. High-speed RICE units commonly reach simple-cycle efficiencies in the 40% to 48% range, well above industrial SCGT and competitive with, or better than, aeroderivative turbines. They achieve this without the capital and schedule burden associated with a combined-cycle plant.

Just as important is how well that efficiency holds up away from full load. Turbine efficiency and output degrade meaningfully below roughly 50% to 70% load. Reciprocating engines hold their efficiency far more consistently across a wide load range, a real advantage for any application where demand varies rather than sitting at a constant, near-full-load setpoint.

Ambient temperature compounds the gap, and elevation compounds it further. As I mentioned earlier, at roughly 2,700 feet in the Permian Basin, a gas turbine is already derated before temperature is even a factor. Add a 115 degree F day and total output can fall 40%. Recovering that lost output on a turbine typically means adding inlet air chillers, and even then, chillers only address the temperature half of the problem; they cannot restore the output lost to elevation.

In contrast, reciprocating engines can be specified with up-graded turbochargers and aftercoolers built for high-ambient, high-altitude conditions that avoid both derates entirely. That capability has to be requested and paid for rather than coming in the default configuration of every unit. Specified correctly, the absence of derating provides a meaningful difference for prime power sized to a summer peak.

RICE also handles load transients better than any gas generation technology compared here. A reciprocating engine's response to a sudden load step, a compressor starting, a drilling rig ramping, or a data center load stepping in blocks, is faster and more stable than a turbine's, as a turbine has to manage rotor thermal stress and compressor surge margins through any rapid change in output. Some engine configurations start from a full stop to full load in single-digit seconds, a characteristic proven over decades of real-world high-speed RICE deployment.

For behind-the-meter applications without a grid to absorb transients, that response speed is often the deciding factor in whether the plant can support the load at all. It is worth noting that many behind-the-meter plants compared in this piece, RICE, SCGT, or aeroderivative, are typically paired with battery storage, a synchronous condenser, or similar equipment to

“

Aeroderivative units arrive largely pre-packaged and factory-assembled, requiring less extensive balance-of-plant work than a small industrial or frame-class turbine typically demands.

”

replace grid functions like ride-through and frequency support. This comparison evaluates generation technology on its own merits rather than the storage or grid-forming equipment layered on top of it.

The value of spares is where RICE's modularity translates most directly into reliability. Gas turbines, industrial or aeroderivative, are strong performers but not immune to downtime: scheduled major inspections every several years, each requiring the unit offline for multiple weeks, combined with forced outages, put realistic turbine availability in the low-to-mid 90% range. For any behind-the-meter application without grid backup during that downtime, that availability gap has to be engineered around, typically with a spare unit.

On a plant built around one or two large turbines, sparing means duplicating a substantial block of capacity. A RICE fleet, built from many smaller engines, can instead carry a reserve margin spread across the whole fleet rather than a conventional N+1 model built around large individual machines. A single engine outage, planned or forced, is absorbed by that margin without any reduction in deliverable capacity, and maintenance can be sequenced unit by unit with no customer load impact.

That same reserve margin approach also supports better heat rate in practice: rather than running every unit in the fleet partially loaded to meet demand, a RICE plant can run only the units needed, at or near full load, where reciprocating engines are most efficient, and hold the rest in reserve. A plant built around a small number of large turbines rarely has that option; meeting reduced demand usually means running the whole machine at partial load instead.

RICE generators' modularity extends to installation and deployment. The units arrive as fully packaged engine-generator

skids, genuinely closer to a plug-in installation than any turbine technology discussed here. Lead times for high-speed units typically run around two to three years, or longer for medium-speed configurations, but they are well inside the multi-year queues facing new gas turbines. Also, units can be commissioned and added to a fleet incrementally rather than as a single all-or-nothing capital commitment.

Where RICE struggles to compete is at the top end of thermal efficiency: a combined-cycle plant, once built, will outperform any RICE fleet on a pure fuel-to-power basis. Power density per unit is lower than a single large turbine block. On cost, RICE typically carries the lowest installed capital cost of the technologies compared here, offset by higher operations and maintenance costs than a turbine of comparable output.

In many applications, the efficiency and availability gains RICE delivers help offset that higher O&M relative to gas turbines. For operators facing variable loads, high ambient temperatures, tight timelines, or high-reliability behind-the-meter requirements, RICE

is frequently the technology of choice across both oil field and data center applications.

Stepping beyond the four technologies compared above, a handful of emerging technologies are drawing attention on the data center side. Solid oxide fuel cells and linear generators offer near-silent, lower-emissions onsite power with few moving parts, and both have found early footholds in hyperscale and enterprise data center projects. The economics remain a real constraint for now: fuel cell installations commonly run at roughly double the installed capital cost per kilowatt of comparable reciprocating engine capacity. Still, the pace of deployment suggests these are technologies worth watching as the market matures. □

“

High-speed reciprocating internal combustion engine units commonly reach simple-cycle efficiencies in the 40% to 48% range, well above industrial SCGT and competitive with, or better than, aeroderivative turbines. They achieve this without the capital and schedule burden associated with a combined-cycle plant.

”



PARKER GREGG

Parker Gregg is vice president of business development and marketing for HiVolt Energy, which provides long-term power solutions to businesses that do not have the luxury of waiting to connect to the grid. Gregg has more than 25 years of progressive experience across the energy and industrial infrastructure markets. Before joining HiVolt in January, he served as principal director of business development at Burns & McDonnell Engineering, a 13,000-person EPC firm, where he spent more than a decade as a senior commercial leader and client sponsor for large, complex energy and power infrastructure programs.

Parker holds a B.S. in mechanical engineering from Texas A&M University. He serves as vice chair of the Texas A&M Mechanical Engineering Industry Advisory Council and is a former board member of the Rice Global Engineering & Construction Forum.