

WHITE PAPER

DIRECT-DRIVE TRANSVERSE FLUX MOTORS FOR DATA CENTER HEAT REJECTION SYSTEMS

A technical overview of large-diameter, low-speed fan drives for data center cooling, and the impact of direct-drive transverse flux fan motors for improved PUE and optimized WUE.

Tyler Williams

Head of Technology and Innovation, ETM

Scott Reynolds


Director of Technology Commercialization, ETM

EXECUTIVE SUMMARY

Data center heat rejection is facing significant technical challenges, driven by increasing rack power density, site water restrictions and electrical grid constraints, as well as tightening noise limits and energy budgets.

Optimizing system design has led to considerable industry debate on the use of evaporative cooling versus closed-loop dry cooling, as well as hybrid and adiabatic options. The optimal system type may vary for each given application, but the thesis of this paper is that in all cases these drivers point to larger, slower-spinning, direct-drive fans, and that Transverse Flux Motor technology (TFM) is the most effective choice to power these fans.

Transverse flux motors provide direct drive capability in the higher-torque / lower-speed operating range required to unlock the beneficial fan physics of large, slow-spinning axial fans. This results in higher efficiency, lower costs and improved reliability, all at one-half to one-third the weight of conventional machines. At a facility level, the primary advantages include significantly lower energy use for improved Power Usage Effectiveness (PUE) and the flexibility to leverage this improved performance to optimize Water Usage Effectiveness (WUE), while also offering lower noise and requiring less maintenance.



The physics driving these benefits are fewer drives per unit of airflow, lower operating pressures, gearbox elimination, and inherent transverse flux low speed efficiency. Each of these mechanisms concentrates torque in the drivetrain, which has historically been the economic constraint on fan and cell consolidation. Gearboxes enable cost-effective AC induction motors but carry reliability and noise penalties, whereas conventional direct-drive motors are heavy and expensive at these torque levels. TFM resolves this tension, being economically competitive with ACI+Gearbox while offering the reliability, noise, and efficiency benefits of direct drive.

A case study is provided to illustrate the potential benefit of TFM for a 1 GW data center. The example projects a 43% reduction in annual air-movement-related heat rejection energy costs, from \$15M to \$8.6M per year, while reducing the required number of fan drives over 70%, from 608 to 175 units. Additionally, the weight of each drive is reduced by more than half due to a significant reduction in raw materials via a smaller electric motor and elimination of the gearbox. See Table 1 "TFM DD 3" vs "Baseline".

BACKGROUND: DATA CENTER COOLING AND HEAT REJECTION TRENDS

As illustrated in Figure 1, industry projections place global data center capacity demand at roughly 170–220 GW by 2030, up from about 80 GW in 2025 [1], with AI clusters operating at 5–10 times the rack density of traditional cloud servers.




Figure 1 – Projected global data center capacity demand through 2030, showing a range of 171–298 GW.

All this IT load is rejected to the atmosphere as heat following four trends:

- Water restrictions.** Evaporative cooling for data center heat rejection consumes one to two liters of water per kilowatt-hour of IT energy [2], tracked as Water Usage Effectiveness (WUE). This equates to millions of liters per day for a typical hyperscale campus. US data centers are projected to consume 150–300 billion liters (40–80 billion gallons) per year by 2028 [3]. Given this reality, closed-loop, zero-evaporation heat rejection system designs are increasingly popular for new hyperscale builds, driven by both ESG (Environmental, Social, and Governance) sustainability commitments and water-use regulations. The most publicized example is Microsoft's Fairwater AI campus, which rejects heat at gigawatt scale through 172 twenty-foot fans in a zero-water, liquid-cooled architecture [4]. However, optimal cooling architecture remains site- and region-specific, and hybrid or adiabatic designs will continue to make sense where water availability and climate allow. The fan efficiency and drivetrain arguments in this paper apply equally in those configurations.
- Energy efficiency.** While shifting from evaporative cooling to dry cooling for heat rejection can benefit on-site WUE, it can increase electrical power consumption, posing significant challenges to Power Usage Effectiveness (PUE). Furthermore, thermoelectric power generation itself consumes significant amounts of water. Without substantial investments to secure wind and solar supply, excess fan energy shifts water consumption upstream to the power producer, offsetting a potentially significant amount of the on-site water savings [5].
- Direct-to-chip liquid cooling (DTC).** DTC for cooling IT equipment in data centers provides much more efficient thermal transfer, which allows higher return-water temperatures from the server. This in turn makes dry air rejection more viable even in extreme heat, expanding the viability of this option in lieu of evaporative cooling to minimize WUE. High-temperature liquid cooling also allows for “free cooling” via dry coolers, effectively eliminating mechanical refrigeration [6].
- Noise limits.** Data center noise has become a leading source of community opposition. Established clusters sit in populated suburbs, latency-sensitive workloads keep new capacity near metro areas, and even rural AI campuses have residential neighbors; noise

ordinances are tightening in all settings. More airflow means more fan and fan drive noise, motivating new approaches.


Taking these trends together, water restrictions motivate a change to increase the use of dry air rejection with larger, slower, higher-torque fans, while noise and energy budgets constrain the design approach. The missing piece has been an efficient, cost-effective, direct-drive motor for this operating point.

To be clear, these same slower, higher-torque fans offer similar advantages in evaporative cooling, hybrid and adiabatic systems, but the dry-cooling transition is accelerating their adoption.

AN INTRODUCTION TO TFM AND AN ANALYSIS OF BENEFIT AT THE FACILITY LEVEL

TFM AND HOW IT DIFFERS FROM OTHER ELECTRIC MOTORS

A transverse flux machine creates a magnetic flux path transverse to the direction of rotation, rather than radially or axially as in conventional machines, as shown in Figure 2. This arrangement decouples the electrical and magnetic circuits, with a single toroidal coil per phase linking every pole. By decoupling the electric circuit from the magnetic circuit, TFM allows pole count to be optimized without proportionally reducing conductor area. This increases continuous torque per size/weight/cost vs. a conventional radial-flux PM motor. These benefits, coupled with the fundamental physics of axial fans, enable the cost-effective direct-drive capability needed to drive large, slow-spinning fans.



AN ILLUSTRATIVE EXAMPLE OF TFM BENEFIT FOR GIGAWATT-SCALE DATA CENTER HEAT REJECTION

Before detailing the underlying aerodynamic and electromagnetic physics, it is useful to establish the macro-level impact of this technology. To contextualize these benefits for a hyperscale facility, the analysis in Table 1 compares illustrative TFM architectures against legacy fan-drive options across various data center heat rejection scenarios.

Table 1. Comparison of 1 GW Data Center Heat Rejection System Scenarios

Parameter	Unit	Baseline: 9 ft geared cell + AC induction motor		TFM DD 1: Direct scale-up	TFM DD 2a: Eased for efficiency and noise	TFM DD 2b: Evap scale up, eased for efficiency and noise	TFM DD 3: Further consolidation	TFM DD 4: PUE Neutral
Cooling	-	Evap.	Dry	Dry	Dry	Evap.	Dry	Dry
Fan Diameter	ft	9	9	14	14	14	24	30
Fan Speed	RPM	439	439	282	250	250	119	72
Tip speed	ft/min	12,400	12,400	12,400	11,000	11,000	9,000	6,800
Airflow per fan	kCFM	197	197	477	423	423	1018	1202
Total Pressure	in. H ₂ O	2.10	2.10	2.10	1.36	1.36	0.91	0.40
Static Pressure	in. H ₂ O	1.51	1.51	1.51	0.88	0.88	0.59	0.22
Velocity Pressure	in. H ₂ O	0.60	0.60	0.60	0.47	0.47	0.32	0.18
Heat Removal Efficiency	kCFM/MW	46	120	120	150	46	178	270
Fan shaft power†	Hp	41	41	96	55	55	87	45
Shaft torque	Nm	662	662	2,417	1,556	1,556	5,171	4,484
Drivetrain	-	ACI+GB	ACI+GB	TFM	TFM	TFM	TFM	TFM
Total Drivetrain Efficiency	%	86.3%	86.3%	92.2%	91.7%	91.7%	92.6%	92.6%
Fans per MCFM‡	-	5.07	5.07	2.10	2.36	2.36	0.98	0.83
Fans per GW rejected	-	234	608	251	355	109	175	225
Cooling Power per GW	MW	8.2	21.4	19.5	15.8	4.9	12.2	8.2
Cooling PUE	-	1.008	1.021	1.019	1.016	1.005	1.012	1.008
Water Consumption per GW rejected	acre-ft/day	40	0	0	0	40	0	0
	households	43,000	0	0	0	43,000	0	0
Cooling energy cost per GW	\$M/GW	5.8	15.0	13.6	11.1	3.4	8.6	5.8
Savings compared to Dry Baseline, per GW rejected	\$M/GW			1.4	4.0	2.4*	6.5	9.3
Land per GW rejected §	Acre	0.9	2.3	2.3	3.2	1.0	4.6	9.3
	% of total	0.35%	0.91%	0.91%	1.28%	0.39%	1.85%	3.71%
Sound pressure at 500 m, per 1GW array ¶	dBA	59.5	63.6	63.6	62.5	57.4	59.8	56.7
Drive system weight per GW rejected	MT	76 (ACI+GB)	198 (ACI+GB)	126 (TFM)	96 (TFM)	35 (TFM)	131 (TFM)	168 (TFM)
				277 (RFM DD)	344 (RFM DD)	105 (RFM DD)	438 (RFM DD)	562 (RFM DD)

Notes: Continuous full-load operation at \$0.08/kWh, normalized per million CFM of duty. † Fan rows follow AMCA 201/301 similarity from the baseline fan, with size-effect allowances of +2% to +3% carried through the power, torque, energy, and cost rows. Gearbox at 95% efficiency; VFD at 97% in all columns. Sound pressure estimated at 500 m for a bank delivering one million CFM; relative values follow AMCA 301, with the absolute anchor estimated by the fan specific-sound-power method. Legacy direct-drive weights interpolated from published ratings. ‡ MCFM per unit heat rejected scaled from the reference (highest-static) design by $(SP_{ref}/SP)^{(n/m)}$. The exponents reflect that exchanger pressure drop falls with face velocity faster ($\Delta P \propto V^m$, $m \approx 1.5-1.8$) than heat transfer ($h \propto V^n$, $n \approx 0.6-0.8$); since $n/m \approx 0.42$, a lower-static design needs more airflow yet still nets lower fan power (at the cost of added exchanger surface). § Cooling yard taken at two times total cell plan area (aisles, clearance, electrical); cell plan scales with the fan. Illustrative; actual values depend on fan selection, coil design, duty modulation, and site limits. ¶ Sound estimation per the fan specific-sound-power method [15].

* Compared to evaporative cooling baseline.

Case Review

Baseline: The baseline models a 1 GW facility using 9-foot fans with geared AC induction drives (440 RPM, 660 N·m) moving 197k CFM at 2.10 in. H₂O total pressure. Under evaporative cooling, this setup uses 230+ fans, spending \$6M/year in energy while consuming 40 acre-feet (13M gallons) of water daily. Switching to a dry-cooling baseline eliminates water use but requires ~600 fan drives (nearly 200 metric tons of drivetrain weight), driving annual energy costs up 2.5x to \$15M. Multiple direct-drive TFM scenarios are evaluated against these baselines to demonstrate how transverse flux motors optimize facility performance metrics.

TFM DD 1: Dry Cooling Direct Scale-Up - To look at the effect of diameter increase in isolation, this scenario scales fan diameter from 9 to 14 feet with no change to tip speed and pressure. Direct-drive TFM technology reduces drive count by nearly 60% (from 600 to 250 units) and cuts annual energy costs from \$15.0M to \$13.6M.

TFM DD 2: Scale-Up Eased for Efficiency and Noise (Dry and Evaporative) - Running the 14-foot fan at a lower speed (250 RPM vs. 282 RPM) trades a modest increase in land usage for reduced acoustic output and improved efficiency.

- **TFM DD 2a (Dry Cooling):** Slowing the 14-foot fans increases the drive count from 251 to 355 units relative to TFM DD 1. However, the increased fan efficiency reduces annual energy costs to \$11.1M—saving \$2.5M annually over TFM DD 1 and \$4.0M over the dry-cooling baseline.
- **TFM DD 2b (Evaporative Cooling):** For sites prioritizing evaporative cooling, replacing 234 9-foot geared fans with 109 14-foot TFM direct-drive units cuts annual energy costs by 41% (from \$5.8M to \$3.4M) and improves cooling PUE from 1.008 to 1.005. Noise drops by 2.1 dBA (59.5 to 57.4 dBA). System reliability improves by eliminating gearboxes entirely, reducing total maintenance points by over 50%.

TFM DD 3: Dry-Cooling Further Consolidation - Consolidating further to 24-foot fans at even lower tip speeds reduces the required drive count to 175 units (50% reduction compared to TFM DD 2a) while further improving efficiency and sound metrics. Annual energy costs drop to \$8.6M, yielding a \$2.5M annual savings over TFM DD 2a and a total energy cost reduction of \$6.5M (43%) while also yielding a 4 dBA reduction in sound compared to the original dry-cooling baseline.


TFM DD 4: Minimize WUE Without Increasing PUE - This dry-cooling scenario matches the low PUE of evaporative cooling while eliminating operational water consumption. Achieving parity requires 225 30-foot fans spinning at 72 RPM, which also yields a meaningful ~7 dBA acoustic reduction compared to the dry-cooling baseline.

Matching evaporative PUE purely through dry air motion requires a fundamental shift in cooling infrastructure. The primary thermodynamic lever to maximize airflow per watt is reducing fan tip speed. However, lower tip speeds reduce available static pressure, decreasing heat removal per CFM and requiring a higher total air volume.


The primary trade-offs for eliminating water without an energy penalty are **capex** due to fan scale (scaling from 9-foot to 30-foot fan diameters) and **footprint**. Heat rejection land area increases from 0.35% to 3.7% of total campus footprint—a minor trade-off for rural or suburban

hyperscale campuses, though a notable consideration for land-constrained urban developments.

Figure 3 illustrates the relationship between tip speed, power draw, and land utilization. Ultimately, it demonstrates that large-diameter, TFM-enabled dry cooling can **eliminate data center water consumption at zero energy penalty**, requiring an increased allocation of site acreage.



Summarizing all the scenarios together, normalized by total heat rejection, the comparison is dramatic, as illustrated in Figure 4.



The **TFM DD 2a** scenario rejects 1 GW of heat with 355 drives and less than half the total drivetrain weight, while saving \$4 million per year in energy cost; **TFM DD 3** is even more striking with about 175 drives saving \$6.5 million a year.

The torque and weight rows in the table indicate why this design direction has been difficult to date. For legacy radial-flux PM motors, with fan torque increasing more than threefold in the direct scale-up case and over sevenfold in **TFM DD 3**, both motor size and cost increase dramatically, and the motor struggles to maintain acceptable efficiency. The TFM solution unlocks the beneficial fan physics while reducing driveline weight and increasing efficiency and reliability. The following sections provide details of the underlying physics of both the fan and the TFM motor, and how these work together to provide the benefits outlined above.

FAN SCALING FUNDAMENTALS

DRY COOLING AIRFLOW REQUIREMENTS

The fundamental metric for atmospheric heat rejection is heat removed per unit mass of air. Evaporative cooling uses water's latent heat of vaporization, carrying away 40 to 60 kJ/kg of air (primarily as water vapor). By contrast, dry coolers rely solely on sensible heat transfer governed by the specific heat capacity of dry air, carrying away only 10 to 15 kJ/kg at realistic operational temperature rises (Figure 5). Eliminating water usage thus reduces thermal energy rejection per kilogram of air by a factor of 3 to 5, requiring a **3 to 5x increase in mass air flow** to cool the same thermal load.

FIGURE 5.

Heat Removed per Kilogram of Air

Evaporative cooling removes 3–5× more heat per kilogram of air than dry (closed-loop) cooling.




Figure 5 – Heat removal per kilogram of air for evaporative and dry (closed-loop) cooling.

FAN AFFINITY LAWS

The fan affinity laws (AMCA 201: $Q \propto N \cdot D^3$, $P_t \propto N^2 \cdot D^2$, $H \propto N^3 \cdot D^5$) describe the fundamental physics of what geometric scaling can and cannot deliver [7]. Scale up fan diameter at constant tip speed and constant total pressure, and airflow and shaft power both grow as D^2 : Airflow per watt is unchanged. The acoustic scaling laws (ANSI/AMCA 301) yield a similar result: sound power per unit airflow is constant with diameter. To a first approximation, a larger fan is not intrinsically a better fan. The benefits of large fans therefore come from four factors not


covered by the affinity laws: the number of drives per unit of airflow, the operating pressure itself, departures from geometric similarity, and the drivetrain. Each is explored in the sections that follow.

FEWER, LARGER FANS

Aerodynamically, a single fan of diameter $2D$ and an array of four fans of diameter D provide equivalent swept area and—to a first-order approximation—identical performance. However, analyzing practical implementation reveals critical non-affinity advantages, as illustrated in Figure 6.

A major driver for deploying fewer, larger fans is economic, mirroring lessons from the wind power industry. Wind turbine power generation scales with the square of rotor diameter (D^2), whereas primary capital and operational expenditures—foundations, towers, electrical interconnections, installation, and maintenance calls—are fixed per unit. Scaling rotor diameter amortizes these fixed per-unit costs across substantially higher energy output, driving diameters upward despite unfavorable torque scaling.

Cooling fan arrays follow the same economic logic. Every fan requires a dedicated motor, variable frequency drive, control interface, wiring, structural support, shroud, and maintenance schedule. These per-unit overhead costs are amortized over a swept area that expands as D^2 , while blade, shroud, and structural costs grow more slowly than swept area. The same thermodynamic and financial scaling principles that push wind turbines to megawatt-scale proportions to minimize levelized cost of energy (\$/kWh) favor larger cooling fans to achieve the lowest levelized cost of air delivery (\$/CFM).



RELIABILITY AND RESILIENCE

The intuition that many small fans are inherently more resilient than a few large ones, a common objection to consolidation, conflates two distinct metrics: **failure event frequency** and **overall cooling capacity availability**.

Mission-critical facilities require guaranteed heat rejection, a system-level property governed by per-unit reliability and design redundancy rather than unit count alone. Cooling arrays routinely incorporate $N + 1$ or $N + 2$ redundancy [18]. The fundamental design question is whether fewer, higher-reliability units can maintain availability while reducing operational maintenance. As illustrated in Figure 7, two primary mechanisms demonstrate that they can:


1. **Elimination of Mechanical Transmissions:** In demanding cooling duty, gearboxes and associated lubrication systems are a leading source of component failure and scheduled maintenance. This mirrors wind energy field data, where gearbox faults account for a disproportionate share of unplanned downtime and operational costs [10],[16]. Transitioning to direct-drive topology eliminates this primary wear mechanism entirely rather than merely attempting to mitigate it [17].
2. **Concentrated Reliability Capital:** A reliability budget scales with total thermal load rather than unit count. Concentrating capital into fewer machines enables engineering specifications that are cost-prohibitive across large fan arrays: bearing systems rated for L_{10} service life exceeding 200,000 hours, sealed or automated lubrication pathways, and integrated vibration/thermal monitoring suites [9].

When a single large direct-drive assembly replaces 4 medium geared units or 8 small direct-drive fans, it can be engineered to a far higher reliability standard. System reliability improves through the combined effect of reduced total failure points and higher individual component MTBF (Mean Time Between Failures).

FIGURE 7.

Reliability Through Fan Consolidation

Fewer, larger direct-drive fans reduce drivetrain components while maintaining redundancy.



VELOCITY AND PRESSURE

The most powerful lever for optimizing fan behavior is shifting along the velocity-pressure axis. Unlike affinity-law scaling, increasing fan diameter for a *constant* airflow lowers the discharge velocity, yielding three compounding benefits:

- **Reduced Kinetic Waste:** Fan work consists of useful heat exchanger pressure drop and wasted exit kinetic energy. Because wasted velocity pressure is proportional to v^2 , doubling the discharge area cuts this waste by a factor of four (Figure 8, left).
- **Optimized Exchanger Thermodynamics:** Exchanger pressure drop falls with face velocity to the 1.5–1.8 power, while heat transfer falls only to the 0.6–0.8 power. This trades a required increase in exchanger surface area for a massive net-positive reduction in fan power.
- **Acoustic Mitigation:** Lower pressure permits lower tip speeds. Because sound power scales with the fifth power of tip speed, the noise reduction from slowing the fan vastly outweighs the acoustic penalty of a larger diameter (Figure 8, right).

In sum, growing the fan faster than the airflow simultaneously reduces pressure, power, and noise.

FIGURE 8.




Figure 8 – Discharge (velocity) pressure loss decreases as fan area increases (constant airflow), and sound power decreases significantly as blade tip speed is reduced (constant diameter). Values are illustrative (ref. AMCA 301).

GEOMETRIC SCALE EFFECTS

Secondary aerodynamic phenomena provide additional advantages at larger scales, further compounding the economic rationale for fan consolidation:

- **Reduced Tip Leakage Fraction:** The physical clearance between the blade tip and fan shroud (constrained by structural deflection and manufacturing limits) does not scale proportionally with fan diameter. Because leakage flow scales with circumference (D) and through-flow scales with swept area (D^2), leakage as a fraction of total flow drops proportionally to $1/D$. This suppresses the tip vortex, mitigating a major source of both aerodynamic loss and acoustic noise.

- **Improved Reynolds Number:** As fan diameter scales, the blade chord Reynolds number increases, which directly reduces profile drag. Classical scale-effect corrections attribute several percentage points of overall efficiency gain to this shift alone.

DRIVETRAIN IMPLICATIONS

Consolidating an array into fewer, larger fans concentrates torque in each drive. At constant tip speed and constant duty per unit of swept area, per-fan power grows with the square of fan diameter, shaft speed falls as $1/D$, and shaft torque grows with the cube of diameter. Motor and gearbox size and cost are determined primarily by torque rather than power, so the drivetrain is the worst-scaling component in the system as fans consolidate. The wind industry encountered the same problem: gearbox reliability became a dominant O&M issue [10], and the direct-drive generators developed in response are among the heaviest and most expensive components on the turbine. The practical limit on cell consolidation has therefore been set less by fan aerodynamics (large fans are catalog items) than by drivetrain cost per torque.

Two drivetrain approaches have historically been available for cooling fans in this torque range:

- **Geared and belted drives** supply torque cheaply from a legacy AC induction motor, but introduce 2–5% transmission loss, oil and seal or belt maintenance, alignment sensitivity, significant noise, and the most common failure modes in cooling tower service. At hundreds-to-thousands of fans per campus, gearbox or belt maintenance becomes a significant O&M cost.
- **Conventional direct drive** eliminates the gearbox, but the physics of traditional radial-flux PMSM motors mean the motors rapidly become prohibitively large and costly, as explained in more detail in the following section.

Published field studies support eliminating the gearbox: in a monitored side-by-side installation, a direct-drive conversion reduced input power by about 12% (the gear train, rated 96% efficient, measured 90.2% as installed) and removed gear oil and shaft alignment from the maintenance schedule [9]. The remaining barrier is the motor itself: published conventional direct-drive PM motors in the 1,400 N·m class run roughly 550 to 700 kg at efficiencies in the low 90s [11] — less efficient and multiple times larger than the geared motors they replace.

NOISE REQUIREMENTS AND SOURCES

Acoustic emissions have emerged as a primary constraint (and political pain point) for data centers. Local noise ordinances typically regulate property-line emissions using A-weighted limits, historically 55 to 65 dBA during the day and 45 to 55 dBA at night. However, newly adopted data-center-specific regulations are notably stricter. Model ordinances now utilized across several eastern U.S. jurisdictions limit residential property-line noise to 45–55 dBA during the day and 40–50 dBA at night [8]. Critically, these updated regulations mandate pre-construction acoustic modeling, post-construction verification, and explicit penalties for tonal and low-frequency noise. These provisions directly target the most common community grievance: a continuous, low-frequency hum that standard broadband A-weighting fails to fully capture. Consequently, the practical design basis for a modern facility is a strict, continuous limit of 50 dBA at the boundary, with zero tolerance for tonal peaks.

Within an air-side heat rejection array, the fan driveline generates noise via three distinct mechanisms:

- **Aerodynamic (Broadband):** Dominated overwhelmingly by blade tip speed, with the tip-leakage vortex acting as a secondary contributor.
- **Gear Mesh (Tonal):** Mechanical transmissions generate distinct, narrow-band tones. Human hearing is acutely sensitive to tonal noise, and modern ordinances explicitly penalize it.
- **Electromagnetic (Tonal):** Motor switching and magnetic frequencies generate tones, though these are negligible in a properly engineered machine.

Field trials confirm that converting a geared cell to direct drive at a constant fan diameter reduces sound pressure levels at the source from 82.3 dBA to 74.4 dBA, while eliminating the gear-mesh tone [9]. Once the mechanical transmission is removed, the remaining acoustic lever is purely aerodynamic. Because sound power scales steeply with tip speed, the ultimate acoustic mitigation strategy aligns perfectly with the thermodynamic solution: expanding the swept area.

TRANSVERSE FLUX MOTOR TECHNOLOGY

For high-torque, low-speed, continuous outdoor duty, the transverse flux motor (TFM) fundamentally changes the economics of heat rejection. The defining advantage of the TFM is its ability to generate several times the continuous torque of a conventional radial-flux motor at a fraction of the size, weight, and cost. By deploying a direct drive TFM, facilities can finally capitalize on the aerodynamic benefits of massive fan consolidation (higher efficiency, lower noise, and higher reliability) without the structural penalties that limit conventional direct-drive motors. The remainder of this section explains how TFM architecture achieves these operational objectives.

HOW TFM WORKS, AND WHERE IT FITS

The transverse flux motor (TFM) derives its name from the orientation of its winding relative to the magnetic airgap. In a conventional radial flux machine (RFM), the winding points toward the airgap and is tightly coupled to the machine's pole pitch. In a TFM, the winding forms a toroidal ring that is largely independent of pole count. This seemingly simple geometric change fundamentally alters how torque-producing capability scales with pole count.

The easiest way to understand the distinction is through the familiar motor sizing relationship:

$$T \propto V_{\text{rotor}} \bar{B} \bar{A}$$

where T is torque, V_{rotor} is active rotor volume, $\bar{B}$ is magnetic loading (airgap flux density), and $\bar{A}$ is electric loading (current per unit circumference). For a given motor size, rotor volume is fixed, while magnetic loading is constrained by the limits of the steel and permanent magnets. Consequently, electric loading becomes the primary lever available to increase continuous torque density.

Low-speed torque production naturally favors a high pole count because increasing the number of poles causes the magnetic energy stored in the airgap to change more rapidly with position. However, in a conventional RFM, each additional pole must share the same fixed circumference.

As pole count increases, the slot area available to each pole becomes smaller, reducing the copper area and current-carrying capability per pole. The gain from additional poles is therefore exactly offset by the reduction in electric loading available to each pole. To first order, simply adding poles does not significantly increase the torque density of an RFM.

The TFM escapes this limitation by decoupling the winding from the pole pitch. A single toroidal winding energizes all poles simultaneously, and its cross-sectional area is determined primarily by the axial length of the machine rather than the pole spacing. Consequently, adding poles does not force a reduction in available copper area. Electric loading therefore increases with pole count rather than remaining approximately constant.

This distinction appears directly in the expressions for electric loading:

$$\bar{A}_{RFM} = J c_x h_c$$

$$\bar{A}_{TFM} = J c_z h_c \frac{Z_p}{\tau_p}$$

Where J is continuous current density [A/mm²], c_x and c_z are the circumferential (in the case of RFM) and axial (in the case of TFM) allotment of space for the coil, and h_c is the radial coil height. All three of these factors are driven by the same physics and tradeoffs and tend to be similar across machines. The distinguishing feature is the factor Z_p/τ_p , or single phase axial length over pole pitch. This term directly multiplies the electric loading of the TFM. Physically, this ratio reflects the fact that winding area can grow with axial phase length *and* grow on a per pole basis as pole pitch is decreased. For practical machine geometries, values of Z_p/τ_p between roughly 3 and 12 are common, producing correspondingly higher electric loading and continuous torque density for the same current density and resistive loss.

The advantage is compounded by winding utilization. When a TFM is formed into a rotational machine, the winding closes on itself to form a continuous copper ring. Nearly all of the conductor is active. In contrast, the end turns of an RFM add resistance, losses, and volume while contributing no torque. This improves the effective copper utilization and further increases the performance advantage of the TFM.

Power, Frequency, and Application Fit

This scaling is not unlimited. As pole count rises, magnetic saturation eventually becomes the dominant constraint. Steel can only support a finite flux density, and beyond this point additional current only contributes leakage flux and additional losses. As a result, TFMs possess their greatest advantage in continuous-duty torque production rather than peak transient torque.

Frequency places a second important limit on machine design. Iron losses increase rapidly with frequency and flux density according to the Steinmetz relationship. For a given cooling system, this establishes a practical operating frequency ceiling. RFMs maximize power within that frequency limit by adopting a low pole count and operating at high speed, using a gearbox to

reduce the output speed to match the application. TFMs take the opposite approach. They employ a high pole count to match the application speed directly and use their torque advantage to eliminate or reduce the mechanical transmission.

This is why TFMs are particularly compelling for low-speed, high-continuous-torque direct-drive applications where gearbox losses, maintenance, noise, and reliability are important considerations.


Motor Constant

While torque is often used to compare machines, continuous-duty applications are ultimately limited by thermal performance. Torque rises linearly with current, while copper losses rise with $I^2 R$. Thus, a more meaningful figure of merit is the motor constant:

$$K_m = \frac{T}{\sqrt{P_{loss}}}$$

Motor constant measures how effectively a machine converts thermal loading into continuous torque. Higher values indicate more torque for a given temperature rise and cooling capability.

FIGURE 9.



The TFM is particularly well suited to maximizing K_m for all of the reasons stated above, resulting in exceptionally high continuous torque per watt of copper loss, making the TFM especially attractive for thermally limited direct-drive systems.

OVERCOMING HISTORICAL TFM LIMITATIONS

While early TFM prototypes demonstrated significant theoretical advantages, commercialization was historically limited by material and electromagnetic challenges [14].

- **Core Material:** Traditional TFMs often relied on soft magnetic composite or powdered-metal cores, which suffer from relatively high hysteresis losses, lower saturation flux density, and limited mechanical strength. Modern, high power TFM architecture implements the required three-dimensional flux paths primarily using conventional non-grain-oriented electrical steel (NGOES), improving magnetic performance, structural durability, manufacturability, and cost.
- **Power Factor and Torque Ripple:** Many early TFM topologies exhibited poor power factor and significant torque ripple. Modern TFM designs achieve power factors comparable to conventional AC Induction motors and PM motors, enabling 1:1 motor-to-inverter sizing, while limiting torque ripple to less than 5% peak-to-peak of rated torque.
- **Mechanical Assembly:** Higher volumes have brought a sustained focus to the mechanical and manufacturing aspects of TFM. Compared to early lab demonstrations, modern TFM machines reduce part count by part combination and leverage careful design for automation. In serial production, TFMs demonstrate touch times in single digit minutes and first pass yield in the high 90s.

Addressing these historical limitations has enabled commercially viable TFM architecture. Internal benchmark testing indicates motor constant (K_m) improvements of approximately 3x to 6x per unit weight and volume relative to state-of-the-art radial-flux permanent magnet machines.

DUTY CYCLE FIT

Fan drives are continuous duty loads sized by sustained torque requirements. This is the duty cycle where the TFM motor constant advantage translates most directly into smaller, lighter, and more efficient machines. Modern transverse flux motors operate at 95% efficiency at the low-speed, high-torque rating points typical of fan duty, with lower coil temperatures at rating and correspondingly longer insulation life. Large direct-drive cooling fans align with the transverse flux machine's strengths: high torque, low speed, continuous duty, severe environment. Ultimately, the operational demands of massive direct-drive cooling fans align perfectly with TFM's fundamental strengths, providing the missing technological link necessary for true, maintenance-free aerodynamic consolidation.

CONCLUSION

The thermodynamics of data center heat rejection unequivocally favor larger, slower, direct-drive fans—but the massive torque they require has long stood as an economic and structural barrier. Transverse flux motors (TFMs) remove this barrier, providing the necessary low-speed torque at half to a third the weight of conventional direct-drive options and with superior electrical efficiency.

Unlocking these fan physics enables a cascade of facility-level advantages. By neutralizing the drivetrain bottleneck, TFMs allow operators to drastically cut energy use (improving PUE), leverage performance margins to eliminate water consumption (optimizing WUE), satisfy strict property-line noise ordinances, and eliminate mechanical transmission maintenance.

REFERENCES

- [1] McKinsey & Company, "AI power: Expanding data center capacity to meet growing demand" (Oct. 2024) and "The cost of compute" (2025): global data center capacity demand projected at 171–219 GW by 2030, from ~82 GW in 2025.
- [2] Environmental and Energy Study Institute, "Data Centers and Water Consumption": average data center WUE $\approx$ 1.9 L/kWh (Green Grid metric).
- [3] A. Shehabi et al., 2024 United States Data Center Energy Usage Report, Lawrence Berkeley National Laboratory, LBNL-2001637 (Dec. 2024).
- [4] Microsoft, "Inside the world's most powerful AI datacenter" (Official Microsoft Blog, Sept. 18, 2025) and "Sustainable by design: Next-generation datacenters consume zero water for cooling" (Dec. 2024).
- [5] SPX Cooling Technologies (Marley), "Cooling Towers Offer Water and Energy Savings as Part of Efficient Data Center Cooling Systems," Expert Opinion Whitepaper series: on-site cooling water consumption compared against water consumed upstream in electric power generation.
- [6] ASHRAE. *AI Data Center Energy Performance Framework: Integrated Design Principles*. Developed by ASHRAE Technical Committee 9.9 (Mission Critical Facilities, Data Centers, Technology Spaces and Electronic Equipment), 2024.
- [7] AMCA Publication 201, Fans and Systems; ANSI/AMCA Standard 301, Methods for Calculating Fan Sound Ratings from Laboratory Test Data.
- [8] PennFuture, Model Sound (Noise) Ordinance for data center facilities, used as a template in several eastern U.S. jurisdictions.
- [9] ABB/Baldor, Direct Drive Cooling Tower System, publication BR411: monitored side-by-side geared vs. direct-drive conversion (input power, as-installed gear train efficiency, A-weighted sound levels).
- [10] U.S. DOE / National Renewable Energy Laboratory, Gearbox Reliability Collaborative (est. 2007): premature wind turbine gearbox failures identified as a leading source of unplanned O&M expense; NREL/TP-500-49133 and related publications.
- [11] Published manufacturer ratings for direct-drive PM cooling fan motors in the 1,400 N·m class (e.g., ABB direct-drive cooling tower motors; Wolong 37 kW / 250 RPM PM motor datasheet).
- [12] W. M. Mordey, "Electric Generator," U.S. Patent 437,501 (Sept. 30, 1890), the first transverse flux machine patent.
- [13] H. Weh and H. May, "Achievable force densities for permanent magnet excited machines in new configurations," Proc. ICEM, Munich, 1986.
- [14] "Transverse Flux Machine — A Review," IEEE Access (2022): low power factor and cogging/torque ripple documented as principal barriers to TFM adoption.
- [15] ASHRAE Handbook — HVAC Applications, Sound and Vibration Control; J. B. Graham, specific sound power level method for fan noise estimation.
- [16] J. Ribrant and L. Bertling, "Survey of failures in wind power systems with focus on Swedish plants, 1997–2005," IEEE Trans. Energy Conversion (2007); P. J. Tavner et al., wind-turbine field reliability surveys (WMEP/LWK databases); and NREL, "Report on Wind Turbine Subsystem Reliability," NREL/TP-5000-59111 (2013): across large fleets the gearbox is consistently the largest contributor to turbine downtime per failure.
- [17] J. Carroll, A. McDonald, and D. McMillan, "Reliability comparison of wind turbines with DFIG and PMG drive trains," IEEE Trans. Energy Conversion 30(2):663–670 (2015): direct-drive (PMG) machines show materially lower drive-train failure intensity than geared (DFIG) configurations.
- [18] Uptime Institute, Tier Standard: Topology (concurrently maintainable and fault-tolerant N+1/N+2 architectures); ASHRAE Datacom Series, Thermal Guidelines and design considerations for data-center cooling: N+1 and higher redundancy is standard practice in mission-critical heat rejection.
- [19] ASHRAE Handbook — Fundamentals, Psychrometrics: sensible heat capacity of moist air, used for the dry-cooling heat-duty and per-gigawatt airflow conversions; energy-cost basis taken at \$0.08/kWh, conservative against the U.S. EIA average industrial retail price of ~8.6¢/kWh in 2025 (Electric Power Monthly).

The performance, energy, cost, weight and noise figures in this paper, including the 1 GW case study, are illustrative estimates based on modeling and the assumptions stated. They are not guarantees of performance. Actual results will vary with site conditions, system design and operating profile.

© 2026 Electric Torque Machines, Inc. All rights reserved.