KULR

NYSE: KULR

KULR TECHNOLOGY GROUP, INC.

Letter to Shareholders

From Michael Mo, CEO and Founder · June 2026

Dear shareholders, customers, and partners,

I want to step outside the cadence of quarterly reporting and share with you, in my own words, where KULR stands today, and where I believe we are headed. We have reached a point where the company's broader strategic vision is coming into focus, and I want to share that vision with you directly.

Battery Is Infrastructure

Let me start with the idea everything else in this letter rests on: battery is infrastructure.

In the digital era, artificial intelligence runs on infrastructure we can see — power lines, fiber, and data centers. With the physical AI era now arriving, intelligence moves into machines that operate out in the world, and those machines depend on a different kind of infrastructure. Every satellite, every drone, every robot, every rack of backup power for AI compute runs on a battery system. The battery is not a component bolted onto physical AI; it is the energy foundation physical AI is built on.

There is a specific reason the battery is the foundation, and it sits at the heart of how we are building KULR. The markets we serve — autonomous platforms, directed-energy systems, and digital infrastructure — look unrelated on the surface, but they share one technical constraint: power density. A drone, a robot, a satellite, a directed-energy system, a rack of AI backup power — none of them needs a battery that simply stores energy. They need a battery that can deliver power: at five to twenty times the discharge rate of a standard cell, sustained through repeated high-demand cycles, with the heat that output generates managed without failure. That is a categorically harder problem than just energy storage, and it is the problem the KULR ONE platform was built to solve.

Power is the wedge. It is why our platform wins design, and everything downstream — the customers, the programs, the revenue — follows from solving it first. The constraint does not relax as physical AI scales; it tightens, with every system demanding more power, in less space, more safely, generation after generation. The company that owns that layer — safe, dense, high-discharge power delivered as a

complete system — owns the infrastructure physical AI runs on. That is what we mean when we say battery is infrastructure: not energy you store, but power you can trust, everywhere the grid does not reach.

That is why our mission for 2026 is as direct as it sounds: build more batteries, and sell more batteries. It is not a slogan — it is the work of laying the infrastructure layer for the systems that will define the next decade. On our last earnings call, I said 2026 would be measured by three things: product revenue growth, gross margin improvement, and cost discipline. That is the commitment, and what we are accountable for delivering. Everything else in this letter is built on top of it.

The first quarter showed real progress: revenue nearly doubled year-over-year, product sales grew sharply, gross margin expanded meaningfully, and operating expenses came down even as revenue grew — the early signs of the operating leverage we promised. We will report the quarter-by-quarter details on our earnings calls; this letter is about the strategy those numbers are building toward. One quarter does not make a turnaround, but the direction is exactly what we said: build more batteries, sell more batteries, operate with discipline. That is the foundation everything else is built on.

What We Are Building On Top of the Foundation

As we develop all the technology pieces for our KULR ONE platform, we are building the energy and power electrification platform for physical AI — the autonomous, mobile, and intelligent systems that operate in the physical world.

Let me explain what I mean by physical AI, because the term is common but its substance is often missed. The AI most investors have encountered lives inside data centers — it runs on GPUs and draws power from the grid. Physical AI is the same intelligence — perception, planning, reasoning — embedded inside systems that operate in the physical world: a satellite processing data in orbit, a robot maintaining a space station, a drone flying an inspection route, a humanoid working in a warehouse, a counter-drone system responding in milliseconds.

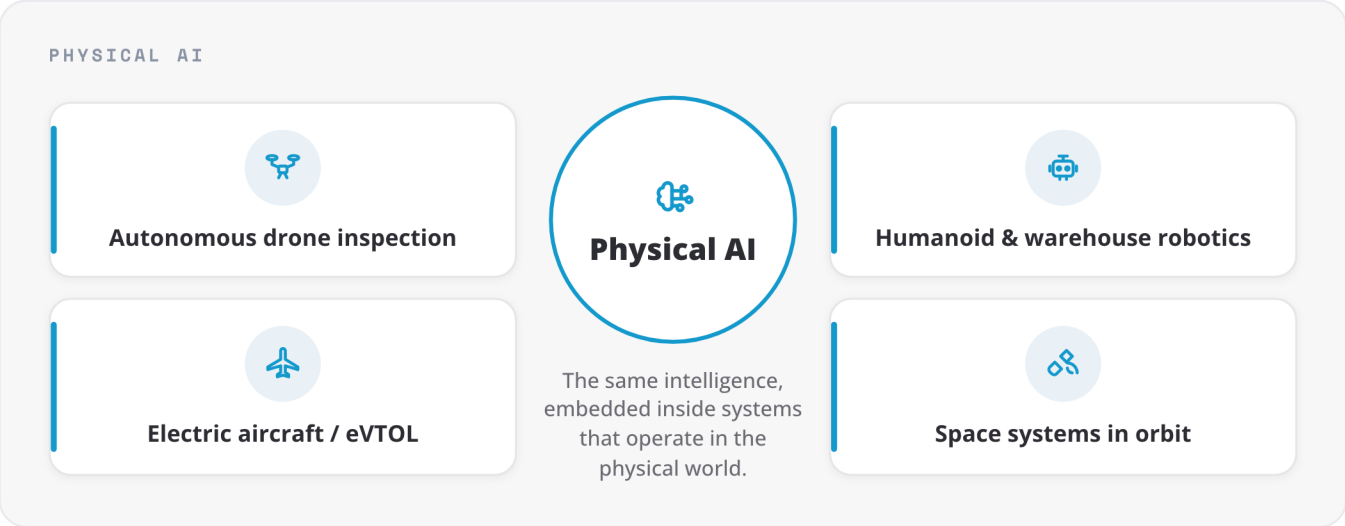

Every one of these systems shares the same constraint. It must carry its own intelligence, because cloud latency makes remote inference impractical and often unsafe. And it must carry its own energy, because there is no grid in the sky, in orbit, on the ocean, on a battlefield, or on a robot floor. Physical AI is therefore defined by the convergence of three disciplines that have historically lived in separate industries: artificial intelligence, energy storage, and power electronics. The companies that integrate across them will define the next decade of physical infrastructure; the companies that operate inside only one will be commodities. This is the structural insight our platform is built on. And the way we get there is to start with what I described: build more batteries, sell more batteries. Every pack we ship is one more proof point that the platform works.

The NVIDIA Lesson

NVIDIA spent more than thirty years building the accelerated computing platform — graphics, then general-purpose parallel compute, then CUDA as a software ecosystem developers could not easily leave — and then watched the world's most important workload, artificial intelligence, land on their architecture as if it had been designed for it all along. The platform was decades in the making; the payoff arrived in a compressed window once the workload showed up. Two lessons sit inside that history, and both shape how we think about KULR.

The first lesson is that platform companies reveal themselves one capability at a time, until the architecture that was always there becomes visible to everyone else. For most of those thirty years, NVIDIA looked like a graphics-card company. It was, in fact, assembling the substrate for modern AI. The second lesson is that platforms compound: each capability reinforces the others and makes the next one easier to add. The value is not in any single component but in the integration — which is what competitors find hardest to replicate. A company selling one component competes on price; a company with an integrated platform competes on architecture, and architecture is durable.

Our mission is to build KULR on a similar path, with one meaningful advantage NVIDIA's own platform has now made possible — an advantage they did not have at the start: artificial intelligence is now accelerating the rate at which platforms can be designed, simulated, qualified, and brought to market. The same intelligence NVIDIA's platform serves is what helps us iterate faster on cell chemistry selection, thermal architecture, control software, power electronics integration and manufacturing design. What took NVIDIA decades, I believe can compress meaningfully — not because the engineering is easier, but because the tools are categorically more powerful than they were even five years ago. To be candid, we are early. What I am committing to is that we will build with the patient discipline that defined the great platform companies, while taking full advantage of the accelerants that did not exist before — and let the architecture reveal itself through what we ship.

KULR's Evolution

If the NVIDIA lesson is about how a platform is built, there is a second lesson — about how a company evolves over time — and the clearest example of it is SpaceX. I raise it because the company KULR is becoming is a natural evolution of the company we have been building: not a pivot, but a progression.

A little over two decades ago, SpaceX began with one hard problem: reaching orbit affordably. It solved that, then made launch reusable, then used that foundation to build Starlink, a global connectivity platform — and today that same orbital infrastructure is being positioned for the AI era, with disclosed plans for constellations of compute satellites in space. SpaceX turned one technology business into the next; it compounded them. Each stage was built on the domain expertise of the one before it, and over roughly twenty-four years a launch company became foundational infrastructure for the next era of computing.

KULR's arc rhymes with that, on our own scale and timeline. The hard problem at our core is older than the company itself: for nearly forty years, the thermal management, carbon fiber, and safety engineering that keep high-energy systems from failing in the most unforgiving environments have been proven in space, alongside NASA and on real space missions. KULR was founded about thirteen years ago to build on that heritage — to carry four decades of space-proven thermal and safety engineering into new applications beyond space and defense. That expertise was never the destination. It was the foundation, because the hard part of building a safe, high-power battery is precisely the thermal and safety engineering that heritage gave us.

That foundation became the KULR ONE battery platform we operate today. The next stage is the same evolution carried forward: from a battery platform into a physical AI energy infrastructure platform — the company that supplies the safe, dense, high-power energy layer that autonomous machines depend on, across every market physical AI is creating. The thermal expertise made the battery platform possible; the battery platform makes the energy infrastructure platform possible. We are not changing what we are — we are growing into the fuller expression of it.

The Platform

Let me describe what the platform actually consists of, because “platform” is easy to claim and harder to substantiate. At the core is the KULR ONE battery architecture — cells and packs engineered for the power density I described, built for high discharge and the thermal stability to sustain that output safely. It is, by design, battery-cell-agnostic: it pairs with whatever chemistry serves the application best, so we can partner with every cell manufacturer and our customers always get the best technology for their needs. As cell chemistry advances and commoditizes through its maturity cycle, the architecture that integrates those cells safely and reliably captures more durable value.

Around that core sit the capabilities that turn a battery into a system: our battery management systems and control electronics; NASA-grade thermal management and passive propagation resistance — the safety engineering that lets a high-power pack operate next to people, processors, or astronauts; and KULR VIBE, our vibration-mitigation technology for the rotors and rotating systems that airborne platforms depend on. We are also beginning to build power electronics organically: the KULR ONE Charger, planned for 2026, will incorporate a power supply unit of our own design — our first power conversion product engineered in-house.

Each piece is useful on its own. Together they form the complete energy and power stack that an autonomous system needs. We are not assembling a catalog of products; we are assembling an integrated platform where the battery, the management software and electronics, and the thermal and safety engineering are designed to work as one.

And we are building the capability to make it at scale. From our vertically integrated facility in Texas — which we are expanding with new high-volume production lines — we are bringing battery assembly, certification, and high-performance component fabrication in-house, so we can build, qualify, and ship faster and at lower cost. A platform is only as real as the factory behind it, and we are building ours to be the one-stop shop the US market needs for high-power batteries.

We bring this platform to five core end markets where physical AI is creating the largest infrastructure opportunities of the next decade.

The Five End Markets

01 · SPACE & DEFENSE

Autonomous systems

The engineering reference standard — KULR ONE Space, qualified in low-Earth and geostationary orbit.

02 · LOW ALTITUDE ECONOMY

Drones & UAS

Below 3,000 feet, toward roughly $210 billion by 2045.

03 · AI DATA CENTER BACKUP

Power at the rack

Edge inference, a roughly $255 billion market by 2030 — on the ground and in orbit.

04 · ENERGY AS A SERVICE

Power delivered as a service

Guaranteed uptime, not equipment — turning hardware sales into recurring revenue.

05 · ROBOTICS

Physical AI on the ground

Toward roughly $370 billion by 2040 — engaged with two humanoid customers; operations in Japan.

The first is space and defense autonomous systems — the engineering reference standards for everything else we build. They operate where battery failure is not recoverable, imposing certification, safety, and reliability requirements no commercial application can match. Meeting that bar in our KULR ONE Space program is what gives our platforms credibility in every other market: customers in defense drones, electric aviation, and AI data centers inherit a battery architecture qualified in low-Earth and geostationary orbits.

That heritage is now extending into physical AI in orbit. Autonomous, free-flying space robots are embodied AI systems that must carry both their own intelligence and their own energy in the most demanding environment that exists — and KULR ONE Space is being selected to power them. Alongside continued satellite mission wins across low-Earth and geostationary orbit, these programs extend our space heritage into a new class of mission. In the most recent quarter, additional low-Earth and geostationary programs selected KULR ONE Space, and our space-qualified batteries remain in active deployment across multiple satellite missions.

The second is the Low Altitude Economy — for a US audience, simply the drone and unmanned aerial systems economy: UAVs and drones operating below 3,000 feet across logistics and last-mile delivery, agricultural and infrastructure inspection, public safety, and the fast-growing fleet of defense and counter-drone platforms procured under NDAA-compliant mandates. Bank of America Global Research projects the global low altitude economy growing toward roughly $210 billion by 2045, and the United States market is opening rapidly as domestic, NDAA-compliant supply becomes a national priority. 2026 is the inflection year — when frameworks become revenue. Every one of these aircraft is, at its core, a battery-powered flying computer, and our KULR ONE Air platform — with a dual-purpose architecture spanning traditional rotorcraft and emerging electric aviation — positions us across this market. Execution here is the furthest along of any market we serve: our high-power flight packs are already in production and broadening adoption, our rotorcraft and electric-aviation partnerships extend the

platform across traditional and emerging aircraft, and we recently won a prototype contract for a US defense drone program — with manufacturing scaling toward thousands of packs per month to meet the demand. And the value of these batteries does not end when their flight life does. A pack engineered for electric aviation retains meaningful useful life once its aviation service is complete, and we are designing for it to begin a second life as stationary energy storage, delivering years of additional service on the ground. One battery, two lives: a more sustainable and more capital-efficient model that turns what the industry treats as end-of-life into the start of a second mission.

The third is AI data center backup — an opportunity spanning two environments converging on the same need. On the ground, AI economics are shifting decisively toward inference at the edge, in telecom facilities, commercial real estate, and distributed sites close to where data is generated — which analysts expect to be the majority of a roughly $255 billion inference market by 2030 (MarketsandMarkets). KULR ONE MAX is engineered for these deployments: high-power, propagation-resistant battery backup that installs at the rack, co-located with compute, without the cooling and footprint of a hyperscale facility. In orbit, the same logic plays out on a larger scale: SpaceX's recent S-1 disclosed plans for up to one million orbital AI compute satellites targeting 100 gigawatts of capacity beginning in 2028 — and because orbital infrastructure cannot be serviced by technicians, it depends on autonomous space robotics for inspection, repair, and assembly, the same systems KULR ONE Space is being selected to power. On the ground or in orbit, AI compute needs energy engineered for power density, safety, and reliability, and KULR is positioned to power both. We are already executing: licensing our propagation-resistant safety and thermal IP to data center OEMs, advancing a high-power backup platform purpose-built for the rack, and holding a seat in the consortium defining next-generation data center power standards.

The fourth is Energy as a Service — mission-critical power delivered as a managed service rather than sold as hardware. If battery is infrastructure, this is how we deliver and monetize it. KULR provides the battery systems, safety architecture, monitoring, and lifecycle management, and the operator pays for guaranteed power, not equipment — turning one-time hardware sales into multi-year recurring revenue and moving backup power off the customer's balance sheet. We are starting where the need is most acute, with telecom operators migrating from lead-acid to lithium-ion — already moving from concept to engagement, with a growing set of operators evaluating the model with us. But the model is not telecom-specific: the same logic of guaranteed uptime, delivered as a service, extends to commercial real estate, data centers, and any infrastructure where downtime is not an option. It is the infrastructure-as-a-service layer of our platform — the same shift that turned computing into a service, applied to power.

The fifth is robotics — physical AI on the ground, and ultimately perhaps the largest opportunity of all. McKinsey projects the general-purpose robotics market growing from under $1 billion in 2025 to roughly $370 billion by 2040; venture funding has tripled since 2023, governments have declared embodied AI a strategic priority, and SoftBank called physical AI its next frontier in its $5.4 billion acquisition of ABB's robotics division — the capital and the conviction are arriving together. Every general-purpose robot faces the same constraint as every other physical AI system: it must carry its own energy, deliver high burst power for dynamic motion, manage heat in a compact enclosure, and stay safe around people and

in a fall. The differentiator is not only the cell chemistry but the pack architecture, thermal management, and operational safety wrapped around it — precisely the KULR ONE platform's strength. Our work here is already underway: through KULR ONE Air we are engaged with two humanoid robotics customers, our space programs already power robotics in the most demanding environment that exists, and we are establishing operations in Japan, one of the world's deepest robotics ecosystems, to position KULR at the energy and safety layer where, as the supply chain matures, durable value will concentrate.

Why Power, Compute, and Intelligence Will Integrate at the Edge

I want to share one structural insight foundational to how investors should think about KULR's place in the future of AI and physical AI. We are not creating that future — it is driven by forces far larger than any one company — but we see clearly where it is heading, and we are positioning KULR to align with this future. As edge AI matures, the relationship between the energy system and the compute system is inverting, and the company that owns the power infrastructure is positioned to own substantially more than power.

Four trends point in this direction. Edge inference silicon is shrinking fast — a Jetson Orin Nano delivers 40 trillion operations per second at 15 watts, smaller than a deck of cards. Small language models are advancing toward distilled forms that run on hardware fitting inside a battery enclosure. Agentic workloads — predictive maintenance, anomaly detection, energy optimization — operate on exactly the current, voltage, temperature, and cycle data the battery management system already holds, making the BMS their natural home. As compute becomes the smaller element, the rational configuration is compute inside the power system, not power beside it — and the owner of the power infrastructure becomes the natural integration point for the compute, memory, and intelligence that run on top of it.

There is a larger architecture implied by all of this. The first era of AI was built on centralization — vast, power-hungry data centers concentrating compute in a few places. Physical AI runs the other way: when intelligence has to live where the work happens — in orbit, in the air, on the factory floor, at the edge of the network — energy and compute must be distributed there too. The future of AI infrastructure is not only larger central data centers but a distributed fabric of energy-and-compute nodes across the physical world. Each of our markets is a node in that fabric where distributed energy and distributed intelligence meet.

I want to be measured about this. It is a structural direction over a multi-year horizon, it will be contested, and it will require KULR to invest in capabilities adjacent to our platform — software, edge AI deployment, and partnerships with model and compute providers. The decisions we are making — the battery-cell-agnostic architecture, the investment in battery management systems, the engineering depth we are extending into Japan — are the decisions that position us at the integration point of the edge intelligence stack as it emerges.

Taken together, the markets this addresses are vast — edge AI inference, general-purpose robotics, the Low Altitude Economy, orbital AI infrastructure, and energy services for critical infrastructure — served by a common platform, the integration of energy, compute, and intelligence at the edge.

Looking Forward

Over the years ahead, we will reveal the platform one capability at a time. Each quarter will bring proof points — customer wins, program advances, manufacturing milestones, partnership extensions, financial discipline — that together demonstrate the architecture we are building. Some quarters will be lumpy, because foundational programs in physical AI are multi-phase and revenue does not always land in the quarter a strategic position is secured. We will be clear about which milestones are foundational and which are revenue-generating.

We will continue to invest in this platform, extend our partnerships, and expand our global footprint with conviction — and operate with discipline, deploying your capital where it builds the most enduring positions. The opportunity, as I see it, is to build the platform the autonomous and intelligent systems of the next decade will depend on — because battery is infrastructure, and that infrastructure is ours to build. The way we get there is by doing exactly what we said we would in 2026: build more batteries, and sell more batteries.

Thank you for the trust you have placed in KULR. I am honored to do this work on your behalf, alongside a team that shows up every day to earn it.

Sincerely,

Michael Mo

Chief Executive Officer and Founder
KULR Technology Group, Inc.

KULR

Market Data Sources

- General-purpose robotics (~$370B by 2040, from <$1B in 2025): McKinsey & Company, “Will embodied AI create robotic coworkers?” (June 2025).
- Low Altitude Economy (~$210B by 2045): Bank of America Institute / BofA Global Research, “The ‘low-altitude’ economy is taking off” (June 2025).
- AI inference market (~$255B by 2030): MarketsandMarkets, AI Inference Market (2025).
- Orbital AI compute satellites (up to ~1 million, ~100 GW from 2028): SpaceX, Form S-1 registration statement filed with the U.S. Securities and Exchange Commission (2026).

Forward-Looking Statements

This letter contains forward-looking statements within the meaning of Section 27A of the Securities Act of 1933 and Section 21E of the Securities Exchange Act of 1934. Forward-looking statements may be identified by words such as “believe,” “expect,” “intend,” “plan,” “will,” “should,” “could,” “may,” “anticipate,” “project,” “target,” “on a [year] horizon,” and similar expressions.

These statements include, but are not limited to, statements regarding the Company’s strategic direction, market opportunities, platform development, partnerships, supply chain, geographic expansion, anticipated benefits of strategic partnerships, anticipated benefits of expansion into Japan, anticipated growth in addressable markets including space and defense, AI inference, AI data center infrastructure, orbital AI infrastructure, the Low Altitude Economy, general-purpose and humanoid robotics, and Energy as a Service for critical infrastructure, anticipated technology roadmap, expected timing of manufacturing capacity expansion and consolidation activities, anticipated future integration of compute, memory, and agentic intelligence with the Company’s power platform, and the Company’s overall business outlook.

Forward-looking statements are based on management’s current expectations and assumptions and are subject to known and unknown risks, uncertainties, and other factors that may cause actual results, performance, or achievements to differ materially from those expressed or implied by the forward-looking statements. Factors that could cause actual results to differ include, but are not limited to: risks related to the Company’s reliance on third parties; risks related to the closing and execution of strategic partnerships and customer agreements; market acceptance and adoption of the Company’s products and services; risks related to the development and certification of new products and platforms; competition; supply chain, geopolitical, and regulatory risks; the timing and execution of manufacturing capacity expansion; risks related to the development of edge AI compute integration and adjacent capabilities; and the other risk factors described in the Company’s filings with the Securities and Exchange Commission, including its most recent Annual Report on Form 10-K and Quarterly Report on Form 10-Q.

The Company undertakes no obligation to update or revise any forward-looking statements to reflect events or circumstances after the date of this letter, except as required by law. Statements concerning third parties, including SpaceX, NVIDIA, Bank of America Global Research, McKinsey & Company, and industry market sizing, are based on publicly available information and are referenced for context. The Company makes no representation as to the accuracy or completeness of such third-party statements.