TESLA THE MAN

AI, the Energy Wall, and a Peek Beyond It

Lev Neymotin, PhD
N5Digital, Inc.
September 2026

We do not need Tesla's laboratory back.
We need the scale of his questions.

INTRODUCTION

WE NEED TESLA BACK. BADLY.

We need Nikola Tesla back. Badly.

Look at the frenzy around AI data centers. What started as a fascinating technical race - bigger models, faster chips, better answers - has spilled out of the computer and entered the physical world. Now we are talking about land, power plants, transmission lines, substations, turbines, transformers, cooling water, neighborhood opposition, utility bills, and politics. The cloud has come down to Earth, and it is much heavier than its name suggested.

Some of the land problem can be handled. Build differently. Build farther away. Build underground where it makes sense. But underground does not make the electricity problem disappear. The electrons must still arrive, and almost every watt that arrives eventually reappears as heat that must leave. We can hide a building. We cannot hide thermodynamics.

The scale of the demand is the real shock. AI is developing on the time scale of software, while the electrical system grows on the time scale of power stations, permits, concrete, copper, and public patience. A new model can appear in months. A new transmission corridor may take a decade. We may not be running out of energy in the universe, but we are beginning to run into a wall in our ability to produce enormous, dependable electric power exactly where and when the new intelligence wants it.

And here Nikola Tesla walks into the story - the genius of electricity, the man with a mind fitting the task.

Not because he would bring us one forgotten invention from a dusty notebook. Not because every grand idea associated with his name worked. Tesla belongs here because he thought in systems. He did not stop at the motor, the generator, the transformer, or the wire. He saw the whole chain: how power is made, changed, transported, controlled, and finally turned into useful work. Near the end of the nineteenth century, that chain changed civilization. In 1900, he even published an essay called *The Problem of Increasing Human Energy*, with special attention to harnessing the Sun. More than a century later, we have manufactured a new kind of intelligence, and the old title almost rewrites itself: the problem of increasing machine intelligence - and feeding it. [1, 2, 3]

If I were putting on Tesla's hat, fitting it on my head as tightly as possible, I would not begin by asking for another power plant. I would ask three more troublesome questions.

First: Why should intelligence need so much power?

Second: If it still needs enormous power, where should we get it?

Third: Must we move the power to the intelligence, or can we move the intelligence closer to the power?

Those three questions open three doors. Behind the first is the low-power human brain, analog computing, neuromorphic systems, synthetic neural materials, and even living neurons placed on chips. Behind the second is the Sun above the atmosphere, where there are no clouds, no night in the ordinary terrestrial sense, and no local argument over whose backyard receives the solar field. Behind the third is Tesla's old dream of wireless power - and then a reversal even he might enjoy: perhaps instead of beaming gigawatts down to terrestrial computers, we should put at least some of the computers up there and beam back only the useful information.

That last possibility is no longer a decorative science-fiction paragraph. I watch SpaceX development carefully, and the line of movement is becoming hard to ignore. Reusable rockets. Mass-produced satellites. Optical links. Large orbital networks. AI joined directly with a launch company. And now a public SpaceX program called Starmind, described as solar-powered AI computation in orbit, with stated design targets reaching 250 kilowatts at peak and 175 kilowatts on average per satellite. The Federal Communications Commission has accepted a SpaceX application for orbital data centers for filing. Acceptance for filing is not approval, and a design target is not an operating machine. But the direction is no longer hidden. [26, 27]

Do I know exactly what Elon Musk has in his mind? Of course not. But I can see the architecture. Launch, solar collection, communications, computation, and mass production are being brought under one roof. That is a Tesla-class move - not a gadget, but an attempt to own the entire physical chain. It intersects directly with the idea that computation may move closer to its energy source.

The temptation is to jump immediately into orbit. I am going to resist it for a moment. Tesla would probably resist it too. Before we race to collect more power, we should ask whether our present machines are wasting power because we chose the wrong computational path.

> *That is where the real story begins.*

THE MACHINE WITH AN ELECTRIC METER

WHEN THE CLOUD BECAME AN INDUSTRY

AI first reached most of us through something almost weightless: a box on a screen. We typed a question and received an answer. The machinery behind it stayed out of sight. That invisibility made the new intelligence feel almost immaterial - language without a body, thought without a metabolism.

But machine intelligence has a very large body. It lives in halls of processors, memory, switches, pumps, fans, chillers, cables, batteries, backup generators, substations, and high-voltage lines. It eats electricity twenty-four hours a day. It produces heat twenty-four hours a day. When demand grows quickly enough, it begins to reshape the places around it.

The numbers have now become too large to remain background scenery. The International Energy Agency estimated global data-center electricity use at roughly 485 terawatt-hours in 2025 and projects about 950 terawatt-hours by 2030 in its central case. That is nearly a doubling in five years. AI-focused data centers are expected to grow even faster. In the United States, Berkeley Lab put data-center use at 176 terawatt-hours in 2023. Its June 2026 update now estimates a 2030 reference case of 649 terawatt-hours - about 11.8 percent of all U.S. electricity - with an uncertainty range stretching from 521 to 843 terawatt-hours. The range is wide because the future is moving too fast for false precision. Even its low end is enormous. [5, 6]

A terawatt-hour is difficult to feel. A gigawatt is easier. Run one gigawatt continuously for a year and it consumes almost nine terawatt-hours. A single giant AI complex measured in gigawatts is therefore not merely a large customer. It is an industrial power system attached to a computing system.

This is why the AI story is becoming political. The issue is not only whether a company can pay for electricity. The issue is whether the electricity can be generated, transmitted, and delivered without taking capacity from other users, raising local costs, or forcing infrastructure decisions on communities that never voted to become part of an AI supply chain.

I do not want to turn this into panic. The electrical system is not collapsing tomorrow. Chips are becoming more efficient. Models are being compressed. Cooling is improving. Software is learning to do more with less. But total demand can keep rising even while each individual task becomes cheaper. We have seen this pattern many times: lower cost opens more uses, more users, and more ambitious applications. Efficiency saves energy per task, then success multiplies the number of tasks.

The AI energy wall is therefore not a single wall standing at a known date. It is a narrowing passage. On one side, intelligence expands quickly. On the other, power infrastructure expands slowly. The danger is that the most consequential technology of our time becomes limited not by ideas, but by watts.

And that leads straight to Tesla's first question, the one that may be more radical than building power stations in orbit:

> *Why should intelligence require so many watts in the first place?*

TESLA'S FIRST QUESTION

WHY DOES INTELLIGENCE NEED SO MUCH POWER?

The obvious answer is that modern AI performs an astonishing amount of computation. That is true, but it is not the end of the answer. It may not even be the most interesting part.

Today's machines are digital. They represent information through discrete states, move it repeatedly between memory and processors, and carry out immense numbers of carefully controlled operations. Digital computing won because it is stable, exact, reproducible, programmable, and manufacturable. A weak zero can be restored into a clean zero. A weak one can be restored into a clean one. Copies can be exact. Errors can be detected. Programs can travel from one machine to another and still behave like themselves.

This was one of the greatest engineering victories in history. I am not declaring war on it.

But the world being represented is not naturally digital. Light, sound, pressure, temperature, movement, chemical concentration, electrical fields, biological signals, and most of what we call perception arrive continuously. We convert that flowing world into numbers because our machines require numbers. Then we make those numbers imitate perception, association, judgment, and language through mountains of digital operations.

It works - magnificently. It also gets hot.

Heat is not merely an inconvenience here. It is a message from the physical world. It tells us that the abstraction has a cost. Every time information is moved, switched, copied, stored, corrected, and moved again, energy is spent. A modern AI accelerator is a miracle of engineering, yet at scale the miracle begins to look like a furnace built to imitate thought.

Then the human brain enters the room and quietly embarrasses us.

THE BRAIN'S QUIET PROVOCATION

The human brain runs on roughly twenty watts. That comparison must be handled honestly. A brain is not executing the same benchmark as a supercomputer, and we do not know how to translate memory, consciousness, perception, language, and movement into one clean number of operations per second. Still, the brain is an existence proof of something enormous: useful, adaptive, general intelligence does not naturally require a private power plant. [7, 8]

The brain does not keep computation in one place and memory in another, shipping every piece of information back and forth across a rigid electronic hierarchy. Its memory is entangled with the structure that processes. Its signals are electrical and chemical. Its activity is massively parallel. It is noisy, approximate, self-adjusting, and distributed. It does not wait for one global clock to command every event.

Biology also computes under merciless energy limits. A brain cannot ask the body for another megawatt. It must compress. It must generalize. It must ignore almost everything and notice what matters. It must act before all information is complete. It survives by using structure rather than brute force.

This is why I have been attracted to analog computing for years. I have always admired the analog world in which we live - the continuous flow of action, graceful and beautiful, before we chop it into digital pieces. My question was never whether we should resurrect room-sized analog machines from the middle of the twentieth century. The real question is whether physical processes themselves can again carry a larger part of computation, now with modern materials, fabrication, sensors, AI calibration, and hybrid digital control. [9]

The future is unlikely to be purely analog. Purity is not the point. Digital should remain where exactness, storage, security, accounting, and error correction matter. Analog should take over where the physics can perform the work more naturally: pattern recognition, vector operations, similarity, optimization, adaptive control, and continuous sensing.

This is not the end of digital computing. It may be the end of digital monopoly.

FROM BITS TOWARD CONCEPTS

There is a deeper possibility. Perhaps the problem is not only that the machine uses digital switches. Perhaps it is also what the machine treats as the basic unit of work.

Digital machines manipulate bits. Software organizes the bits into numbers, tokens, vectors, images, models, and eventually what we recognize as meaning. The physical machine itself remains far below that meaning, switching and moving tiny fragments with tremendous discipline.

Human thought appears to work at a different level. We do not rebuild the idea of a tree from individual sensory atoms every time we see one. We operate through patterns, relationships, compressed meanings, associations, and whole conceptual shapes. Present AI already uses vectors, embeddings, latent spaces, and attention - mathematical structures in which meaning is partly represented by position and relationship. But all of that is still simulated through digital machinery underneath.

This led me to Neural Vector Architecture, or NVA: not one chip and not one product, but a direction. The machine would move some of its work from bits toward continuous vectors, stable patterns, attractors, and concept-like structures. Digital systems would supervise, store, verify, and communicate. Analog cores would perform the operations that analog physics can perform directly. The measure of progress would not be transistor count alone. It would be useful concept density - how much meaningful structure can be represented and transformed per unit of energy. [10]

This is still a proposal, not an accomplished architecture. “Concept” is a dangerous word in engineering unless it becomes measurable. What is a concept physically? A distributed voltage pattern? An optical interference state? A stable region in a dynamic material? A vector stored across many analog elements? The answer may be different for different tasks.

That uncertainty does not weaken the question. It defines the research.

And once we ask what a concept might look like in matter, the chip itself begins to change.

MAKE THE MATERIAL DO THE WORK

A conventional computer is built from components instructed to behave with extraordinary precision. What if instead we design a material whose natural response already contains memory, delay, threshold, relaxation, and adaptation?

That is the idea behind physical reservoir computing. Feed a changing signal into a richly dynamic physical system. Let the system's own internal evolution spread that signal through many interacting states. Then use a relatively simple trained readout to interpret the result. The reservoir is not asked to imitate a standard processor. Its physics is recruited as the computation. [12, 13, 14]

My N5 Synthetic Ionic/Electrochemical Reservoir, N5-SIER, follows that path in a deliberately biomimetic but non-biological way. It imagines a synthetic neural medium in which electronic conduction, ionic movement, electrochemical state, and slower material changes coexist. Fast pathways respond immediately. Ions move more slowly and leave a fading memory. Still slower changes condition the medium. The computation is not placed on top of the material; it emerges from the way the material behaves. [11, 15]

The key lesson from biology is not “put a tiny artificial neuron here, another one there, and connect them exactly like a diagram.” The deeper lesson is to organize matter so that sensing, memory, transformation, and adaptation happen in the same physical neighborhood.

This could lead to low-power processors for temporal signals, sensor fusion, medical devices, robotics, anomaly detection, and adaptive control. It is not ready to replace a data-center GPU. It may never be the right tool for exact arithmetic or general software. But that is not the test. The test is whether it can remove whole classes of energy-hungry digital work by letting matter do what matter already knows how to do.

Optics offers another branch. A lens performs a transformation simply by bending light. Photonic systems can mix signals and perform matrix-like operations at remarkable speed. Recent demonstrations of analog optical and photonic reservoir machines show that physical computation can reach the front edge of modern research, not merely the museum of computing history. Light is excellent at propagation and parallel mixing; electrochemical matter is attractive for memory and adaptation. The interesting future may combine them. [12, 13]

Here the story could remain comfortably synthetic. But biology has opened a more unsettling door.

IMITATE BIOLOGY - OR USE IT?

Researchers are now placing living neurons on electronic interfaces, stimulating them, reading their responses, and using software to create feedback. The DishBrain experiment showed cultured neurons adapting within a simplified Pong-like environment. Organoid-intelligence researchers are exploring three-dimensional neural cultures as possible biological computing systems. FinalSpark has offered remote access to a “wetware” research platform. Cortical Labs markets the CL1 as a biological computing platform in which living neurons interact with silicon. [16, 17, 18, 19]

This is not a replacement for the GPU sitting around the corner. Living neural systems are difficult to grow, maintain, standardize, program, reproduce, and scale. Their useful lifetime, memory, training method, interfaces, and total energy cost - including life support - remain open questions. Even the language becomes treacherous. Adaptation is not the same as intelligence. A neural response is not automatically thought. Company descriptions are not independent proof of general computing capability.

Still, the direction matters. We have moved from studying the brain, to imitating its architecture, to building synthetic materials that borrow its electrochemical style, and now to asking whether living neural matter itself can become part of a machine.

That brings a new engineering constraint that silicon never raised: welfare. As neural cultures become more complex, we will have to ask not only whether they compute, but whether they can experience anything that places moral limits on how they are used. The ethics cannot be attached at the end like a warning label. It grows with the machine.

I find this frontier exciting and uncomfortable at the same time. That is usually a sign that the frontier is real.

QUANTUM: ANOTHER DOOR, NOT THE SAME DOOR

Quantum computing also returns computation to physics. Instead of forcing everything through ordinary binary states, it uses superposition, interference, and entanglement for specialized classes of problems. Google's Willow work demonstrated an important advance in quantum error correction below a key threshold - a real step toward machines whose logical reliability can improve as more physical qubits are added. [20]

But quantum computing is not a general low-power replacement for AI data centers. Its hardware is fragile and highly specialized. Many systems need extreme cryogenic infrastructure. Its advantages appear in selected problems, not in every workload. Quantum may eventually transform chemistry, materials, cryptography, and certain simulations. It may help discover the materials and architectures needed for a new energy system. It does not presently remove the broad AI power wall.

So Tesla's first question does not produce one answer. It produces a family of attacks: better digital algorithms, analog and mixed-signal processors, in-memory computing, photonics, synthetic neural media, biological computing, and quantum tools where they truly fit.

Suppose one or several of these paths work brilliantly. Suppose the energy required for useful intelligence falls by a factor of ten, one hundred, even one thousand.

Are we saved?

Not necessarily.

> *Because cheap intelligence may be used everywhere.*

THE TRAP HIDDEN INSIDE EFFICIENCY

WHEN SAVING ENERGY CREATES MORE DEMAND

Imagine that a future analog or hybrid processor performs an AI task using one hundredth of today's energy. That is a spectacular victory. The data center can shrink. Cooling can shrink. The same grid connection can support much more intelligence.

Then intelligence becomes cheaper.

It moves into every vehicle, laboratory, factory, hospital, home, robot, instrument, and personal device. Models run continuously instead of occasionally. Scientific systems explore millions of designs instead of thousands. Simulations become richer. Machines watch, predict, optimize, translate, diagnose, tutor, and create without stopping. One hundred times more efficiency can invite one thousand times more use.

This is not a reason to reject efficiency. It is a reason not to confuse energy per task with total energy. Efficiency is one side of the solution, not the whole solution.

The human brain is again instructive. Its low power did not cause nature to stop at one brain. Nature built billions of them.

If AI becomes essential to scientific progress, medicine, engineering, education, and management of complex civilization, humanity will not voluntarily freeze it at today's scale. Nor should it. The purpose is not to ration intelligence. The purpose is to make its growth physically sustainable.

So after Tesla has forced us to reduce the hunger of the machine, he asks the second question:

> *Where can we find the power for the intelligence we will still want?*

TESLA'S SECOND QUESTION

WHERE DO WE GET THE POWER?

The first answer is not romantic. It is Earth.

For the next decade and probably much longer, AI will be powered mainly by terrestrial systems. Existing grids must be strengthened. Transmission must expand. Generation must grow. Nuclear power, gas, geothermal energy, terrestrial solar and wind, hydroelectricity, storage, demand management, and better use of waste heat will all compete and cooperate in different places.

There is no honest space solution that arrives fast enough to excuse us from this work. Space-based power does not remove the need to build substations on Earth, manufacture equipment, permit lines, or make difficult choices. NASA's own assessment of space-based solar power described major capability gaps and found representative 2050 systems more expensive than terrestrial alternatives under the assumptions it studied. That is not a death sentence. It is a reminder that orbital power must earn its place against a moving terrestrial target. [21]

But Earth has limitations that do not disappear. Weather. Night. Land competition. Geography. Long transmission routes. Local politics. Fuel logistics. Water. Public fear. A civilization that wants vastly more electricity may eventually need to extend its energy system beyond the surface.

Then the Sun becomes impossible to ignore.

THE SUN WITHOUT WEATHER

At Earth's distance, sunlight arrives above the atmosphere with about 1,361 watts of power per square meter facing the Sun. On the ground, clouds, night, season, latitude, and atmosphere interrupt it. In orbit, the resource is cleaner and more continuous, although not perfectly continuous; eclipses, orientation, degradation, and thermal limits still matter. [25]

The numbers are both encouraging and humbling. With optimistic but reasonable future conversion efficiencies, a continuously delivered gigawatt would require active collection surfaces measured in several square kilometers before adding structure, spacing, storage, maintenance, degradation, transmission loss, and redundancy. That is enormous. It is also not absurd on the scale of orbital industry.

We should correct one instinct from my original outline. We do not want to go “deep into space” to find stronger sunlight. Move away from the Sun and the light weakens. The first sensible energy territory is near Earth: high orbits for fixed relationships with receivers, or other orbital arrangements that trade continuity against distance, beam control, maintenance, and communications.

Europe's SOLARIS studies, Japanese space-solar research, NASA assessments, and Caltech's Space Solar Power Demonstrator all show that the subject has moved beyond pure fantasy. Caltech's MAPLE experiment wirelessly transferred small amounts of power in space and produced a detectable signal on Earth. The amount was tiny. The importance was not the amount; it was the closing of a first experimental loop. [22, 23, 24]

Now the engineering grows savage. Kilometer-scale structures must be launched or manufactured, unfolded or assembled, aimed, cooled, repaired, and kept stable. Solar cells degrade under radiation. Micrometeoroids arrive without appointments. The system must remain economical even while terrestrial energy technology continues improving.

Still, the Sun in space offers something no fuel train can deliver: a vast primary energy flow already crossing the orbit of Earth every second.

Capturing it creates electricity.

And immediately creates the next problem.

There is no cable from orbit to Earth.

> *Now Tesla's old dream returns.*

THE OLD TESLA DREAM RETURNS

POWER WITHOUT WIRES - AT PLANETARY SCALE

Assume we have collected sunlight in orbit and converted it into electricity. The electricity is not yet useful to a data center on Earth. We must change it into something that can cross space, enter the atmosphere, arrive at a receiver, become electricity again, and join the terrestrial grid.

In my original notes I allowed the whole electromagnetic spectrum to enter the imagination - radio, light, X-rays, gamma rays, perhaps something not invented yet. That is the right way to begin dreaming and the wrong way to finish engineering. First principles quickly narrow the field.

X-rays and gamma rays carry high-energy photons, but that does not make them attractive power carriers. They are ionizing, difficult to generate and capture efficiently at utility scale, and unacceptable as giant directed beams through the atmosphere. Higher photon energy is not automatically better engineering.

The serious candidates are microwaves and laser light.

Microwaves pass through much of the atmosphere and can be converted to electricity by a rectifying antenna, or rectenna. They tolerate clouds better than lasers and have a long history of wireless-power experiments. Their disadvantage is physical size. Long wavelengths spread. To keep a beam controlled across tens of thousands of kilometers, the transmitting aperture and receiving field become enormous.

Laser beams can be tighter because the wavelength is far shorter. The transmitter and receiver may be smaller. But clouds, aerosols, weather, atmospheric turbulence, eye safety, and conversion losses become much more severe. A laser-power system may fit selected routes, high-altitude receivers, the Moon, or spacecraft better than dependable grid delivery through ordinary weather.

Neither beam is a magic pipe. Every conversion loses something: sunlight to electricity, electricity to microwave or laser energy, transmission through space and atmosphere, capture at the receiver, conversion back to grid-quality power. A beautiful efficiency at one stage can be destroyed by the chain around it.

THE BEAM THAT REFUSES TO STAY NARROW

Diffraction is the quiet law that ruins many casual sketches. A beam does not travel forever as a perfectly straight cylinder. It spreads. The longer the wavelength, the smaller the transmitter, and the farther the distance, the more it spreads.

At microwave frequencies often studied for space solar power, even a transmitting array roughly a kilometer across in geostationary orbit would illuminate an area on Earth several kilometers wide. That is why serious concepts use very large transmitting antennas and very large rectennas. JAXA's illustrative gigawatt-scale studies speak naturally in kilometers, not in rooftop dishes. [24]

This scale may actually help safety. The beam can be spread over a large area so that the power density is far below the intensity of a weapon. The rectenna can be a sparse mesh, allowing some sunlight and land use underneath. But low intensity demands large area. Physics collects payment one way or another.

Beam control must also be extraordinarily dependable. A utility power transmitter cannot wander across cities, aircraft, satellites, or neighboring countries. It needs pilot signals, automatic shutdown, sidelobe control, cybersecurity, independent monitoring, and international rules. A power beam is not simply a private machine in the sky. It is public infrastructure crossing shared space.

THE LAST MILE IS THE ATMOSPHERE

Space itself is not empty in every practical sense, but for electromagnetic transmission it can be remarkably clean. The difficult last mile is Earth's atmosphere.

Microwaves face absorption, weather effects that vary with frequency, and the need to coexist with communications and radar. Lasers face clouds and turbulence directly. Both face the final conversion into stable electricity. Both need receivers located where the beam can arrive safely and where transmission lines can carry the power onward.

We could place receivers in deserts, offshore, or in regions already committed to energy production. We might distribute many smaller beams rather than depend on one colossal source. We might route power between satellites before choosing the clearest path to Earth. We might invent phased arrays that self-correct with an accuracy impossible today.

Or perhaps we are still thinking in the wrong direction.

Why insist on moving the energy all the way down to the computer?

Why not move the computer up to the energy?

> *That reversal changes everything.*

TURN THE ARROW AROUND

MOVE INTELLIGENCE TO THE POWER

The usual architecture is simple: collect energy wherever it is available, deliver electricity to a terrestrial data center, and keep all computation on Earth.

Turn the arrow around.

Put solar collectors and computers together in orbit. Use sunlight locally. Run selected AI workloads there. Send requests, data, model updates, and results between Earth and space. Instead of beaming continuous gigawatts through the atmosphere, transmit information.

A useful answer may contain far less energy than the computation that produced it. A spacecraft might observe Earth continuously, process enormous streams of raw sensor data in orbit, and send down only anomalies, maps, warnings, or compressed conclusions. Scientific instruments could analyze what they see before transmission. AI inference could be placed near orbital communications networks. Some batch computations might be scheduled when power and thermal conditions are favorable.

This is the brilliant attraction: move bits instead of watts.

But not every workload belongs in orbit. A tightly coupled frontier-model training cluster moves tremendous amounts of data among processors with extremely low latency. Break that cluster into satellites separated by moving distances and the communication system may become the computer's new bottleneck. Tasks that are independent, loosely coupled, delay-tolerant, or already connected to space-based data are far better early candidates. Recent workload studies make exactly this point: orbital compute should begin where location creates a natural advantage, not where it merely reproduces a terrestrial data center at much higher cost. [30, 32]

The first orbital AI systems may therefore be less glamorous than an artificial general intelligence circling Earth. They may filter imagery, manage satellite constellations, perform navigation, monitor weather and climate, detect fires or military activity, compress scientific data, or serve specialized inference close to an orbital network.

That is how large transitions often begin - not by replacing the whole old world, but by taking the tasks that fit the new world naturally.

MUSK'S TESLA-CLASS MOVE

This is where SpaceX becomes central rather than decorative.

I have watched SpaceX build capability in layers. First launch. Then reuse. Then frequent launch. Then a huge satellite network. Then laser links among satellites. Then Starship, intended to move very large mass to orbit. Then the merger of launch, communications, and AI under a connected corporate structure. Now Starmind publicly places solar-powered AI computation in orbit. [26, 28]

Seen separately, each project has its own business. Seen together, they form a ladder.

A company attempting orbital computation needs cheap and frequent launch, standardized spacecraft, solar arrays, thermal systems, optical communications, fault-tolerant computing, autonomous operations, and a reason to deploy at scale. SpaceX does not possess every solution. No one does. But it has assembled more of the chain than any ordinary data-center company could.

That is why I call the move Tesla-class. The comparison is not about personality, mythology, or whether Musk “is the new Tesla.” It is about system boundaries. A normal company asks how to improve its component. A Tesla-class mind asks how many components must be brought together before the entire limitation moves.

Starmind's public design target - up to 250 kilowatts peak and 175 kilowatts average compute power per satellite - is already far beyond the tiny computing budgets of traditional spacecraft. It is also far below a terrestrial hyperscale data center. That middle scale is useful. It forces real engineering without pretending that orbit will absorb the world's AI load next year. [26]

The FCC filing matters for the same reason. It is not proof that the system will work. It shows that orbital data centers have crossed from private imagination into an official deployment pathway. [27]

And then space presents its invoice.

SPACE IS NOT A FREE REFRIGERATOR

People often imagine space as cold. It is not a giant cold bath waiting to cool a computer. Space is mostly vacuum. There is no air to carry heat away and no river to run through a cooling tower. Heat can move through the spacecraft and finally leave mainly as thermal radiation.

Every watt used by a processor becomes almost a watt of heat. A 175-kilowatt orbital computer must therefore radiate roughly 175 kilowatts, plus additional heat from power electronics, batteries, communications, sunlight absorbed by the structure, and other equipment.

At a moderate radiator temperature around 350 kelvin - about 77 degrees Celsius - ideal physics already points to more than two hundred square meters of effective radiating surface for 175 kilowatts. Real hardware needs more because radiators do not see perfect black space in every direction, emissivity is imperfect, geometry matters, fluid loops and pumps consume power, sunlight and Earth add heat, surfaces degrade, and redundancy is not optional. [26, 29]

Raise the temperature and the radiator can shrink dramatically, because thermal radiation rises very quickly with temperature. But processors, memory, lubricants, electronics, and structural materials may not tolerate the temperature. Lower the temperature for ordinary chips and the radiator grows.

This creates a direct marriage between analog efficiency and orbital computation. Cut the computing power in half and you do not merely save half the solar array. You also reduce power conditioning, batteries, internal wiring, coolant flow, and radiator area. In space, every watt removed from the computation removes a chain of mass around it.

A recent engineering preprint examining a one-megawatt orbital data-center concept found photovoltaic and radiator areas measured in thousands of square meters. The exact result depends on assumptions and the work is preliminary, but the message is solid: an orbital data center is a power-and-thermal spacecraft first, a computer second. [30]

Tesla would enjoy that reversal. The problem that looked like computation becomes heat. The problem that looked like energy becomes structure. The problem that looked like structure becomes launch.

RADIATION, REPAIR, BANDWIDTH, AND DEBRIS

Terrestrial data centers replace failed parts constantly. Technicians walk through aisles. Networks reroute traffic. Cooling equipment is serviced. Orbit turns every repair into a logistical event.

Radiation damages electronics and solar cells. Charged particles can flip bits or permanently degrade devices. Shielding adds mass. Redundancy adds mass. Repair robots add mass and complexity. Sending replacements may be easier than repairing, but then the economics depends on launch cost and the responsible disposal of failed spacecraft.

Bandwidth also matters. If an orbital computer requires enormous terrestrial data streams, the communications system may consume much of the advantage. The best workloads will either originate in space, produce compact results, tolerate delay, or exploit an orbital network already in place.

Then there is debris. Thousands of large power-and-compute satellites cannot be treated as disposable glitter. They need collision avoidance, end-of-life control, deorbit or graveyard plans, and international coordination. A brilliant private architecture can become a public orbital burden if this is ignored.

None of these problems invalidates the idea. They define it.

The difference between science fiction and engineering is not lack of imagination. It is the willingness to follow every attractive idea until all of its ugly support systems are visible.

> *Now the two halves of the story can finally see each other.*

THE TWO-FRONT ATTACK

MORE POWER - AND LESS HUNGER

We began with what looked like a simple shortage: AI wants more electricity than the grid can provide quickly.

Tesla's first question attacked the demand. Make intelligence less energy-hungry through better digital design, analog computation, in-memory architectures, photonics, synthetic neural media, biological insight, and specialized quantum tools.

Tesla's second question attacked the supply. Expand terrestrial power now, then explore the Sun above the atmosphere as orbital industry matures.

Tesla's third question attacked the connection between them. Beam power to Earth where it makes sense. Move selected computation to orbit where information is easier to transmit than energy.

The answers multiply each other.

A more efficient processor needs a smaller power plant. In orbit it also needs a smaller solar array, radiator, battery, and launch vehicle. A better power system allows experimental computing substrates to grow without immediately colliding with grid limits. An orbital network can process space-born data locally. Terrestrial analog accelerators can reduce demand while space power develops. Wireless-power research can support the Moon, spacecraft, remote installations, and eventually Earth.

This is not one giant bet. It is a connected portfolio of technologies with different time scales.

The sequence almost writes itself. First, squeeze waste out of chips, algorithms, cooling, grids, and terrestrial generation. As analog and mixed-signal machines begin proving real advantages, move them into the tasks where they fit. In parallel, let orbital computing grow around workloads born in space, push wireless-power demonstrations beyond laboratory scale, and learn how to assemble and service large structures above Earth. Only after those pieces begin closing do space solar, orbital computation, and terrestrial grids become one useful system rather than three beautiful diagrams.

The adventure stays real only if we count the whole machine: energy per useful result, including cooling and conversion; complete orbital mass, not just processor mass; demonstrated performance, not a design target dressed as an operating system. That does not make the vision smaller. It keeps the vision alive long enough to become real.

THE HUNGRY MACHINE BECOMES THE INVENTOR

There is a wonderful circularity here. AI is creating the demand that forces us to redesign computing and energy. The same AI may help design the answer.

The search spaces are enormous. Which analog material composition produces the right balance of memory, noise, stability, and speed? Which phased-array geometry minimizes sidelobes? Which orbit, collector orientation, and thermal schedule give the best delivered computation per kilogram? Which radiator shape unfolds reliably? Which solar cell survives radiation? Which workload should remain on Earth and which should move upward?

Human engineers can formulate these questions. AI can explore millions of combinations, discover patterns in simulation and experiment, control self-calibrating analog machines, and operate orbital systems too complex for continuous human management.

The ultimate hungry consumer may become the tool that invents its own more efficient brain and helps build its own energy supply.

That thought is exciting enough without leaving known physics.

But my original outline did leave it.

It asked whether we might eventually stop collecting energy from stars and learn to touch energy hidden in space itself.

> *Only now, after doing the hard engineering, have we earned the right to open that final door.*

THE LAST DOOR

ENERGY DISSOLVED IN SPACE ITSELF

The Sun is not the only energy in the universe. Space contains electromagnetic fields, particles, gravitational structure, quantum fields, and whatever physical reality lies behind dark matter and dark energy. It is tempting to imagine an ocean of hidden power surrounding us everywhere: no mines, no fuel, no giant orbital collectors - just pick it up and convert it wherever needed.

That would be the ultimate answer to the transmission problem. No cable from space. No microwave beam. No last mile through the atmosphere. The source would already be here.

Today, this remains beyond engineering.

Dark energy is the name we give to whatever is driving the accelerated expansion of the universe. We do not know what it is. The fact that it dominates the large-scale energy accounting of the cosmos does not mean it can be concentrated, tapped, or converted into electricity. Its apparent energy density near us is extraordinarily small, and no mechanism for harvesting it is known. [33]

Quantum fields are also real, and their lowest-energy states are not simple emptiness. But the existence of vacuum fluctuations is not the same as a demonstrated fuel source. Physics has produced subtle protocols involving quantum energy transfer, boundary effects, and temporary local changes, yet no machine has shown net continuous power extracted from the vacuum ground state. [34]

So I will keep the idea where it belongs: not in a development schedule, but at the edge of the map.

That edge matters. Today's established physics was once outside the engineering map. Electromagnetic waves existed before Hertz demonstrated them and before Tesla put them to work. Nuclear energy existed before anyone could control a chain reaction. Quantum mechanics existed in nature before humans found the language to describe it.

The proper attitude is neither ridicule nor announcement. Ask. Search. Refuse to confuse mystery with a power plant.

For now we are stuck - happily stuck - with sunlight, fields, matter, thermodynamics, and the technologies we can begin building.

> *That is already enough for a remarkable future.*

A PEEK BEYOND THE WALL

TESLA THE MAN

Humanity has begun to manufacture intelligence outside the human brain. That intelligence may become one of the greatest multipliers in history. It may accelerate science, discover medicines, design materials, manage complex systems, teach, translate, create, and help us see patterns no unaided mind can hold.

And it arrives with an electric meter attached.

The easy response is to build larger data centers and more power stations. We will do some of that. The dangerous response is to assume that scaling the present architecture is the whole future.

That is why we need Tesla back.

Not to solve the problem for us. Not to hand us a lost wireless-energy secret. Not to replace careful engineering with the romance of genius.

We need the Tesla habit of enlarging the question until the hidden architecture becomes visible.

Why should intelligence need so much power? The question takes us from digital switching to analog flow, from separated memory and processing to materials that remember while they compute, from artificial neurons to synthetic neural media, and from imitating biology to cautiously placing biology itself on chips.

Where can expanding intelligence find power? Earth must carry the immediate load. Then the Sun above the atmosphere appears - immense, continuous enough to be transformative, and difficult enough to occupy generations of engineers.

How should power move? Tesla's dream of wireless transmission returns, now with phased arrays, microwaves, lasers, kilometer-scale apertures, automatic control, and a hard atmospheric last mile.

Then the arrow reverses. Why move all the energy down? Move some of the intelligence up. SpaceX's Starmind program has already placed that question into a public engineering trajectory. Musk's brilliant move is not merely “a computer in space.” It is the attempt to join launch, power, communications, computation, and scale into one system.

And physics smiles, then presents the bill: giant solar arrays, giant radiators, radiation damage, bandwidth, maintenance, debris, and economics. Every door opens into another room.

That is not failure. That is the story of technology.

The future I see is not one machine replacing another. It is a layered intelligence-and-energy system. Digital where exactness rules. Analog where physical continuity can do the work. Biological principles where evolution has already solved energy limits we barely understand. Quantum where the problem truly belongs to quantum mechanics. Terrestrial power for the urgent decades. Space solar and orbital computing as industry moves beyond Earth. Wireless energy where the complete chain becomes safe and economical. Unknown physics kept on the horizon - neither dismissed nor sold before it exists.

Is this a prediction? No. It is a peek beyond the wall, from where we stand now.

Some branches will fail. Some will remain specialized. Some will merge. One or two may change civilization. The purpose of the essay is not to choose the winner before the race begins. It is to see that the race is much larger than “build another data center.”

The machine that consumes the power may help design the new computation. The new computation may make orbital power practical. Orbital industry may make wireless energy practical. The entire chain may become one of the first great projects designed jointly by human and artificial intelligence.

So let us put on Tesla's hat. Fit it as tightly as we can. Remove the restraints of existing technology, but not the restraints of physics. Dream freely, calculate honestly, and let every difficult answer lead to the next, larger question.

Tesla does not return as a man.

He returns as a way of thinking.

NOTES AND SOURCES

[1] Nikola Tesla, "The Problem of Increasing Human Energy, with Special Reference to the Harnessing of the Sun's Energy," *The Century Magazine*, June 1900. Source
[2] Tesla Science Center at Wardenclyffe, historical materials on Wardenclyffe and Tesla's wireless world system. Source
[3] U.S. Energy Information Administration, historical overview of Nikola Tesla's polyphase alternating-current system and its role in generation, transformation, transmission, and motors. Source
[4] Lev Neymotin, "Tesla the Man," original working outline supplied for this essay.
[5] International Energy Agency, *Key Questions on Energy and AI*, Executive Summary, 2026. Source
[6] Sarah Josephine Smith et al., *United States Data Center Energy Usage Report: 2025 Update*, Lawrence Berkeley National Laboratory, published June 2026. Source
[7] Vijay Balasubramanian, "Brain power," *Proceedings of the National Academy of Sciences* 118, no. 32 (2021). Source
[8] William B. Levy et al., "Communication consumes 35 times more energy than computation in the human cortex, but both costs are needed to predict synapse number," *PNAS* 118, no. 18 (2021). Source
[9] Lev Neymotin, "Analog Computing," unpublished working paper supplied for this essay.
[10] Lev Neymotin, *The End of the Digital Age: Toward Analog Intelligence - NVA, Neural Vector Architecture*, unpublished position paper supplied for this essay.
[11] Lev Neymotin, *N5 Synthetic Ionic/Electrochemical Reservoir: A Biomimetic Synthetic Neural Medium Computer*, unpublished concept paper supplied for this essay.
[12] Kirill P. Kalinin et al., "Analog optical computer for AI inference and combinatorial optimization," *Nature* 645 (2025): 354-361. Source
[13] Dongliang Wang et al., "Ultrafast silicon photonic reservoir computing engine delivering over 200 TOPS," *Nature Communications* 15 (2024): 10841. Source
[14] Xiangpeng Liang et al., "Physical reservoir computing with emerging electronics," *Nature Electronics* 7 (2024): 193-206. Source
[15] T. M. Kamsma et al., "Brain-inspired computing with fluidic iontronic nanochannels," *PNAS* 121, no. 18 (2024): e2320242121. Source
[16] Brett J. Kagan et al., "In vitro neurons learn and exhibit sentience when embodied in a simulated game-world," *Neuron* 110, no. 23 (2022): 3952-3969. Source
[17] Lena Smirnova et al., "Organoid intelligence (OI): the new frontier in biocomputing and intelligence-in-a-dish," *Frontiers in Science* 1 (2023): 1017235. Source
[18] Fred D. Jordan et al., "Open and remotely accessible Neuroplatform for research in wetware computing," *Frontiers in Artificial Intelligence* 7 (2024): 1376042. Source
[19] Cortical Labs, CL1 biological-computing platform and company technical descriptions. Company performance characterizations are treated as company claims. Source
[20] Google Quantum AI and collaborators, "Quantum error correction below the surface code threshold," *Nature* 638 (2025): 920-926. Source
[21] NASA Office of Technology, Policy, and Strategy, space-based solar power assessment, January 2024. Source
[22] European Space Agency, SOLARIS and space-based solar-power research. Source
[23] California Institute of Technology, Space Solar Power Demonstrator / MAPLE wireless-power results. Source
[24] Japan Aerospace Exploration Agency, research on microwave wireless-power transmission and space solar power systems. Source
[25] NASA Jet Propulsion Laboratory, "Calculating Solar Power in Space," including total solar irradiance near Earth's orbit. Source
[26] SpaceX, Starmind AI Satellite public program page and stated design targets. Source
[27] Federal Communications Commission, "Space Bureau Accepts for Filing SpaceX's Application for Orbital Data Centers," February 4, 2026. Source
[28] SpaceX, Starship vehicle page and public design description. Source
[29] European Space Agency Advanced Concepts Team, spacecraft thermal-control discussion in vacuum. Source
[30] Slava G. Turyshev, "Orbital Data Centers: Spacecraft Constraints and Economic Viability," arXiv:2604.27197 (2026). Preprint; not peer reviewed at the time of this essay. Source
[31] "Space-CIM: Enabling Compute-In-Memory Accelerators for Thermally-Constrained Space Platforms," arXiv:2606.05741 (2026). Preprint. Source
[32] "Which Workloads Belong in Orbit? A Workload-Centric Framework for Space-Based Compute," arXiv:2603.20317 (2026). Preprint. Source
[33] NASA Science, "What Is Dark Energy?" Source
[34] G. Jordan Maclay, "Gedanken experiments with Casimir forces and vacuum energy," *Physical Review A* 82, 032106 (2010). The paper illustrates why exchanges involving vacuum energy do not supply arbitrary net energy. Source