The Energy Hidden in Information: What the Next Technology Wave Will Actually Be

A physics-grounded forecast of the next technology and market niche, 2026-2036

The question, stated honestly

A tempting idea circulates whenever people notice how much electricity computing consumes: energy is never destroyed, only transformed, so if we pour energy into creating information, surely we can get some other form of energy - or even useful work - back out of that information. If we could, we would have a new energy source and we would stop "losing" energy in conversions.

This article takes that intuition seriously, follows it all the way down to the physics, and shows that it points somewhere real - but not where it first seems. The honest version of the question turns out to be far more valuable than the original.

First principle: information is physical, and a bit can be traded for work

The intuition is not mysticism. It rests on one of the most beautiful results in twentieth-century physics.

In 1929 Leo Szilard imagined an engine driven by a single bit of information. In 1961 Rolf Landauer, working at IBM, made it precise: manipulating information has an unavoidable thermodynamic price. Erasing one bit of information must dissipate at least **kT ln2** of heat into the environment, where k is Boltzmann's constant and T is temperature. Run the same logic in reverse, and one bit of information can in principle be converted into at most **kT ln2** of useful work extracted from a heat bath.

At room temperature (300 K) that number is:

kT ln2 = 2.87 x 10^-21 joules per bit ≈ 0.018 electron-volts ≈ 3 zeptojoules.

This is not a theory waiting for confirmation. It has been measured. Experimental "information engines" - a single electron in a box operated as a Szilard engine, a colloidal particle in an optical double-well trap, an electron level in a semiconductor quantum dot - have all extracted work approaching kT ln2 per bit, with the best runs reaching roughly 0.75 to 0.9 of the theoretical maximum. The bookkeeping of the second law balances exactly. Information genuinely has thermodynamic weight: bits can be traded for work, and erasing them costs heat. This part of the intuition is correct and settled.

So the original question has a real answer: **yes, you can get work directly from information.** The catch is in the amount.

The wall: why "information as an energy source" cannot work

Here is the number that ends the fantasy and opens the real opportunity.

Humanity created roughly 181 zettabytes of data in 2025 - about 1.45 x 10^24 bits. Suppose we could build a perfect Szilard engine and convert *every one of those bits* into work at the full Landauer limit. The total work available would be:

1.45 x 10^24 bits x 2.87 x 10^-21 J/bit ≈ 4,160 joules.

Four thousand joules. That is the kinetic energy of a heavy dumbbell dropped from a shelf. It is the entire annual data output of civilization, converted with perfect efficiency, and it would not boil a cup of tea.

Meanwhile, the data centres that processed and stored that information consumed about **415 terawatt-hours** in 2024 - roughly 1.49 x 10^18 joules. The gap between what we spend on information and the theoretical work-content of that information is a factor of about **3.6 x 10^14**. Three hundred trillion to one.

The conclusion is unavoidable and worth stating plainly: **information is not fuel.** The kT ln2 that a bit carries is real but microscopic. You cannot power a home, a car, or a city from stored data, no matter how much of it you accumulate. Any product, pitch, or patent premised on "extracting usable energy from information at scale" is fighting physics and will lose.

If we had stopped at the original question, this would be a disappointing article. But the same equation, read from the other side, points at one of the largest technology opportunities of the decade.

A word on the "information has mass" hypothesis

A reader following this field will have met a bolder claim, and it deserves a direct answer rather than silence. Since 2019 the physicist Melvin Vopson has proposed a mass-energy-information equivalence principle: that a bit of information carries a small but real rest mass, about **3.19 x 10^-38 kg** at room temperature, so that a full hard drive would weigh microscopically more than an empty one. It is an intriguing idea and has a proposed test (via electron-positron annihilation), but two things must be said clearly.

First, it is unconfirmed - a conjecture awaiting an experiment, not established physics. Second, and more decisively for us, even if it were true it changes nothing about the energy argument. Convert that proposed mass to energy with E = mc-squared: 3.19 x 10^-38 kg x (3 x 10^8 m/s)-squared = **2.87 x 10^-21 J** - the exact same kT ln2 we already had. The mass hypothesis, true or false, does not hide a larger energy reservoir inside information. The wall stands. So we treat it honestly: interesting, unproven, and irrelevant to the business case either way.

Turning the question over: the number that matters is not what a bit contains, but what a bit costs

Landauer's limit has two faces.

Read one way, it says: a bit contains at most kT ln2 of extractable work. That face is a dead end, as we just saw.

Read the other way, it says: **erasing a bit need cost no more than kT ln2 of wasted heat.** It is a floor, not a ceiling - the minimum energy a computation must dissipate. And here is the decisive fact:

today's silicon chips operate a factor of roughly **one million times above** that floor. A modern transistor dissipates on the order of 10^3 to 10^6 times kT ln2 every time it erases a bit.

That million-fold gap is the real "lost energy" the original intuition was groping toward. It is not stored inside our data waiting to be harvested. It is being thrown away, continuously, as heat, every time a conventional processor overwrites a value and discards the old one. The energy problem of computing is not that we fail to extract energy from information. It is that we destroy information wastefully, and the second law charges us for every destruction.

So the genuinely valuable question - the one that survives contact with physics and is stable for the next fifty years - is this:

> **How do we compute and process information while returning as much of the spent energy as possible back into the system, so that each successive bit costs progressively closer to the thermodynamic floor of kT ln2?**

This is not a question about energy generation. It is a question about **energy recovery in computation.** And unlike the first version, it has a clear engineering path, a measurable target (the six orders of magnitude of headroom), and a market pulling on it hard.

Why the market is pulling now

Three facts make this the right decade, not a distant curiosity.

First, demand. Global data centre electricity use is projected to roughly double from 415 TWh in 2024 to about **945 TWh by 2030**, and to reach roughly 1,200 TWh by 2035. At an illustrative industrial electricity price of 0.20 euro per kWh, that 2030 figure represents on the order of **190 billion euro per year** in electricity alone. Grids are straining: in several advanced economies, a fifth of planned data-centre projects risk delay simply because power cannot be delivered fast enough.

Second, the old engine has stalled. For decades the industry could ignore exotic architectures because conventional single-core performance improved around 50% per year. That figure has now fallen below 2% per year. When the free lunch ends, alternative architectures stop being academic.

Third, energy - not transistor count - has become the binding constraint on AI. Inference, not training, now dominates AI energy use and is projected to account for roughly three-quarters of AI energy demand by 2030. The bottleneck the industry is hitting is thermodynamic, and thermodynamic problems invite thermodynamic solutions.

The prediction: which technology and niche appears next

Reading the physics and the market together, the next wave is not "energy from information." It is **near-thermodynamic-limit computing** - hardware that recovers, rather than dissipates, the energy of computation. Two concrete branches are already emerging from the lab into the market, and a third is a natural extension.

Branch 1: Reversible computing (the strongest near-term bet)

If you never erase a bit - if every operation preserves enough information to run it backward - then Landauer's floor does not apply, and in principle the energy can be recycled instead of dissipated.

Practically, this is done with adiabatic, resonant circuits that use inductors and capacitors to slosh charge back and forth, recovering most of it each cycle rather than dumping it to ground.

This has left the whiteboard. A UK-based company, Vaire Computing, taped out a reversible test chip in 2025 whose resonator recovered roughly 50% of the energy in the circuit, with prototypes demonstrating 40-70% recovery, and is targeting a commercial release around 2027. Industry veterans who spent decades on this in national labs now argue that within roughly a decade every new chip leaving a foundry could be reversible. The theoretical prize quoted for the approach is efficiency gains on the order of thousands of times.

A caveat that a serious builder must hold: energy recovered *inside* a resonator can still be lost in the external power supply, and every time a gate genuinely overwrites its output, Landauer still charges for that erasure. Reversible computing pushes the cost per operation toward the floor; it does not abolish the floor. The engineering fight is real and unfinished. But the direction is sound and the momentum is visible.

Branch 2: Thermodynamic / probabilistic computing (the higher-risk, higher-ceiling bet)

Instead of fighting thermal noise, use it. A large class of the most valuable modern workloads - generative AI sampling, Bayesian inference, Monte Carlo optimization - is fundamentally about drawing samples from probability distributions. Conventional hardware burns enormous energy forcing deterministic transistors to *simulate* randomness. A thermodynamic chip lets physical noise do the sampling directly.

This, too, is now silicon. One startup (Extropic) has built probabilistic-bit circuits from standard transistors and published claims of up to 10,000x lower energy than a GPU for the sampling operations they target, with a development platform already beta-tested and a commercial-scale chip planned. Another (Normal Computing) taped out what it calls the first thermodynamic computing chip, aimed at linear algebra and stochastic sampling, with a scaling roadmap through the late 2020s. The 10,000x figures are vendor claims on favorable workloads and should be treated as directional, not audited - but the physical logic (stop paying to fake the randomness you can get for free) is sound.

Branch 3: The white space - an idea to develop, stated honestly

Here is where the two branches suggest something less crowded, and where an entrant could look for original ground - provided we are honest about what already exists.

Reversible computing recovers the energy of *logic*. Thermodynamic computing exploits the physics of *sampling*. And in AI inference, moving data to and from memory often dominates the energy budget, not the arithmetic itself. So the instinct is to point energy-recovery at the data-movement layer rather than the compute core, co-designed with a probabilistic core so the sampler and its memory share one energy-recovery clock.

Honesty requires flagging that the neighbouring ground is already occupied. Adiabatic logic-in-memory architectures apply charge-recovery clocking inside memory arrays. Probabilistic-memory proposals offload the sampling operation itself into the memory fabric. In-memory stochastic sampling on magnetic-tunnel-junction and ferroelectric devices treats device noise as a native random source. Each of these is real and recent. So "recover energy in memory" and "sample in memory" are, separately, taken.

What I could not find as a single, named thesis is the *combination*: reversible-adiabatic energy recovery applied specifically to the **interconnect and data-movement layer** - the wires and refresh, not the array logic - clocked in lockstep with a noise-native sampler so the recovered energy of moving a datum is the same energy that drives the next sample. That precise co-design, centred on transit rather than storage or logic, is where the white space appears to still be. It is speculative, it depends on the reversible-logic and probabilistic-hardware work maturing first, and adjacent teams are clearly circling the same territory. But it sits exactly on the line the physics is drawing, and it is a sharper, more defensible claim than "energy from information" ever was.

What to do with this if you are a builder

The physics gives a compass that will not spin for fifty years, because kT ln2 is not going to be repealed:

Build **toward** the Landauer floor, never claim to break it. Anything advertising net energy *harvested* from information, or computation below kT ln2 per irreversible bit, is a red flag - walk away or short it.

The defensible thesis is **euro per useful operation**, driven down by energy recovery, on the workloads that actually dominate the bill: AI inference, sampling, and the data movement around them.

The near-term, lower-variance play is reversible/adiabatic logic and clocking. The higher-ceiling, higher-variance play is probabilistic/thermodynamic sampling hardware. The unclaimed original play is energy-recovery applied to the memory-and-interconnect fabric, co-designed with a noise-native core.

Information was never going to be a fuel tank. It is a ledger, and the second law audits every entry. The opportunity of the next decade is not to cash the ledger for energy that was never there - it is to stop overpaying on every line.

Sources grounding the key figures: Landauer's kT ln2 and its experimental confirmation (single-electron, colloidal-particle, and quantum-dot information engines); IEA "Energy and AI" projections for data-centre electricity (415 TWh in 2024, ~945 TWh by 2030); public reporting on Vaire Computing (reversible chip tape-out, 40-70% energy recovery), Extropic and Normal Computing (probabilistic / thermodynamic chips). All physical constants and derived quantities computed directly: kT ln2 at 300 K = 2.87 x 10^-21 J; total 2025 data output ≈ 1.45 x 10^24 bits; perfect-engine work-content of that data ≈ 4,160 J; ratio to 2024 data-centre energy ≈ 3.6 x 10^14.