ENERGY EATS INTELLIGENCE

Nuclear Fission

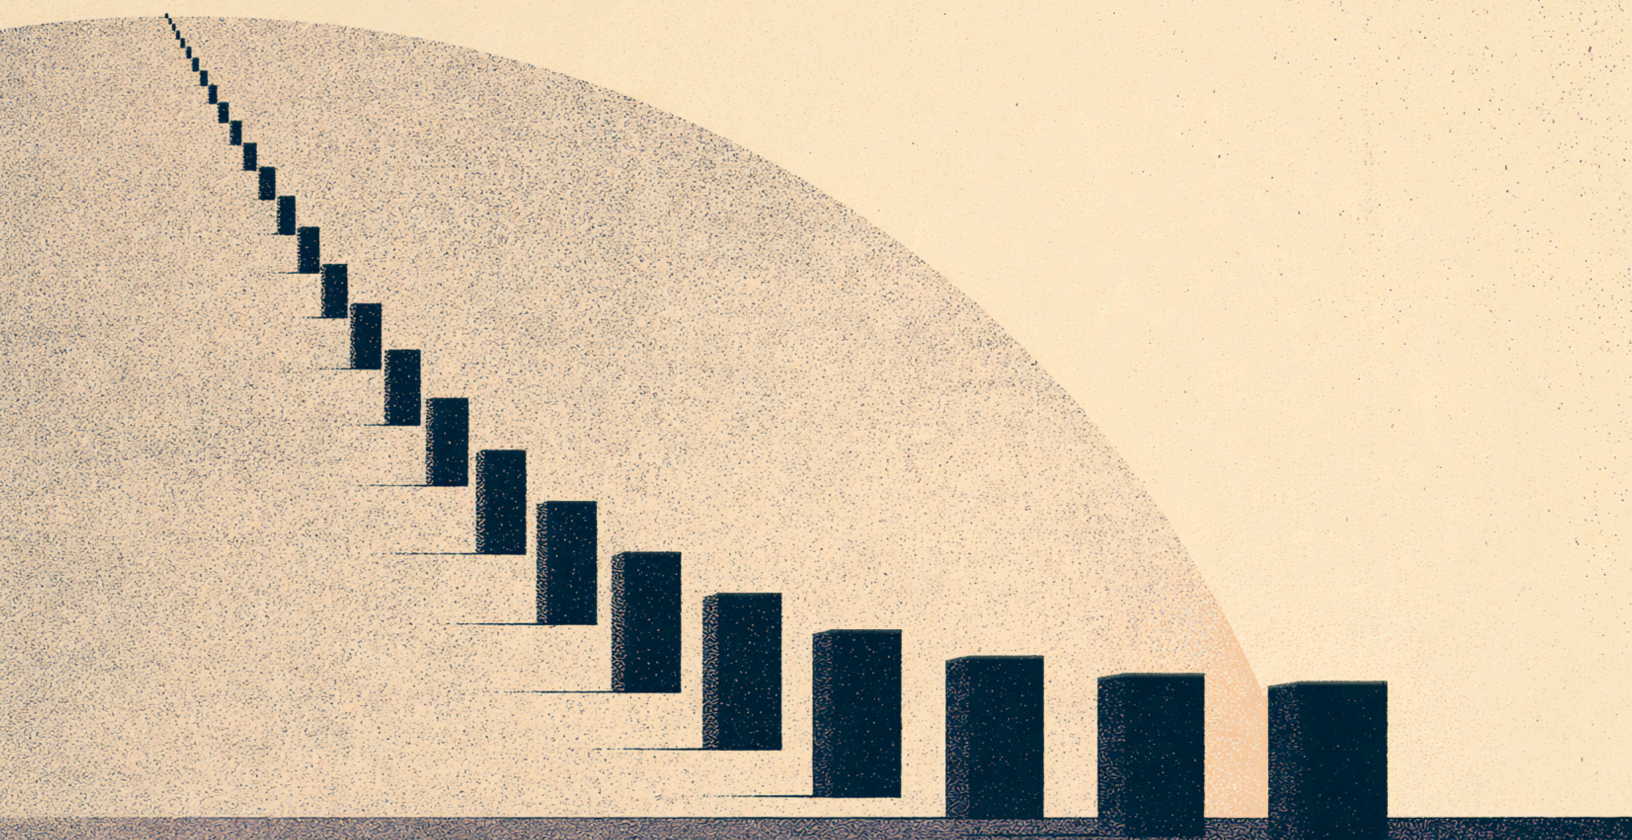

Prime Movers Lab

Roger Carter · July 2026

The reactor was never the expensive part; the megaproject and the decade of financing it were. The bet now winning capital: small, factory-built reactors, funded by anchor customers, not states.

The United States spent about $36 billion and the better part of two decades finishing two reactors at Plant Vogtle in Georgia. In the same span, China connected more than two dozen to its grid on schedule, at roughly $2,500 to $3,000 per kilowatt against Vogtle's roughly $16,000. The American pair came in at 2.6 times their budget and seven years late. Same physics, same proven design, utilizing a technology about as old as the jet engine.

Most people read about Vogtle and concluded that nuclear power is hopelessly expensive, or that the West can't build big things anymore. Neither is a law of nature. What blew up at Vogtle was the cost of building a one-off project on site, not the reactor inside it. This piece is the second in a series; the first argued that AI's buildout is about to leave the United States short of firm power, and surveyed the power-generation options by geography. This one goes deeper on nuclear fission, and on why a wave of capital is now betting that small, factory-built reactors could be a key component to solving the AI energy crisis.

1. It was never the physics

The $36 billion spent at Vogtle is a construction number, the price of a megaproject rather than of the reactor's physics. Look at every large reactor the West has finished this century: France's Flamanville, Finland's Olkiluoto, the UK's Hinkley Point, America's Vogtle. Two reactor designs were reused across these projects. These weren't experiments; they were meant to be the Nth copies of proven machines, and they all still ran two to seven times over budget.

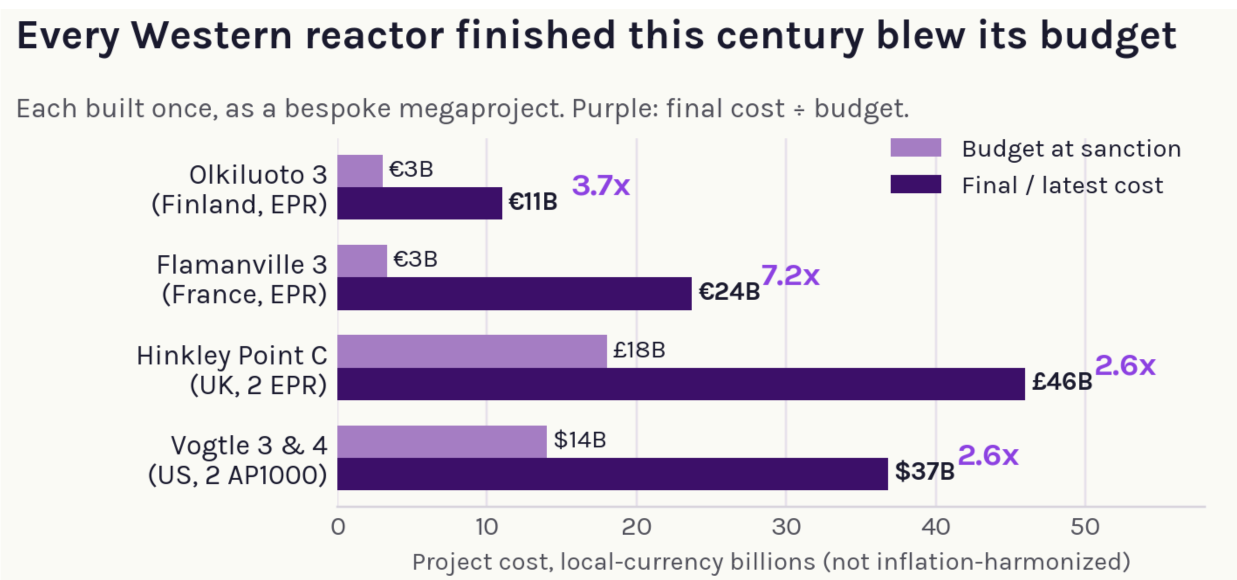

The popular explanation, after Grubler (2010), is that nuclear suffers "negative learning": a technology so complex that costs rise with experience instead of falling. It's a tidy story, and per recent research, it's mostly wrong. Lovering, Yip, and Nordhaus, working through 349 reactors across seven countries, found no such law. Cost escalation was real but specific. In the United States, mostly after the Three Mile Island accident in 1979, a wave of post-hoc regulation landed on plants already half-built, inflating their costs. France escalated mildly and then went flat. However, during the same period, South Korea's costs *fell* by about half over the life of its program, because it froze a standardized design and built it over and over, with a standing workforce and a stable regulator.

Cost escalation was never a law of nuclear

Construction cost diverged sharply by country. The US is the outlier — not the rule.

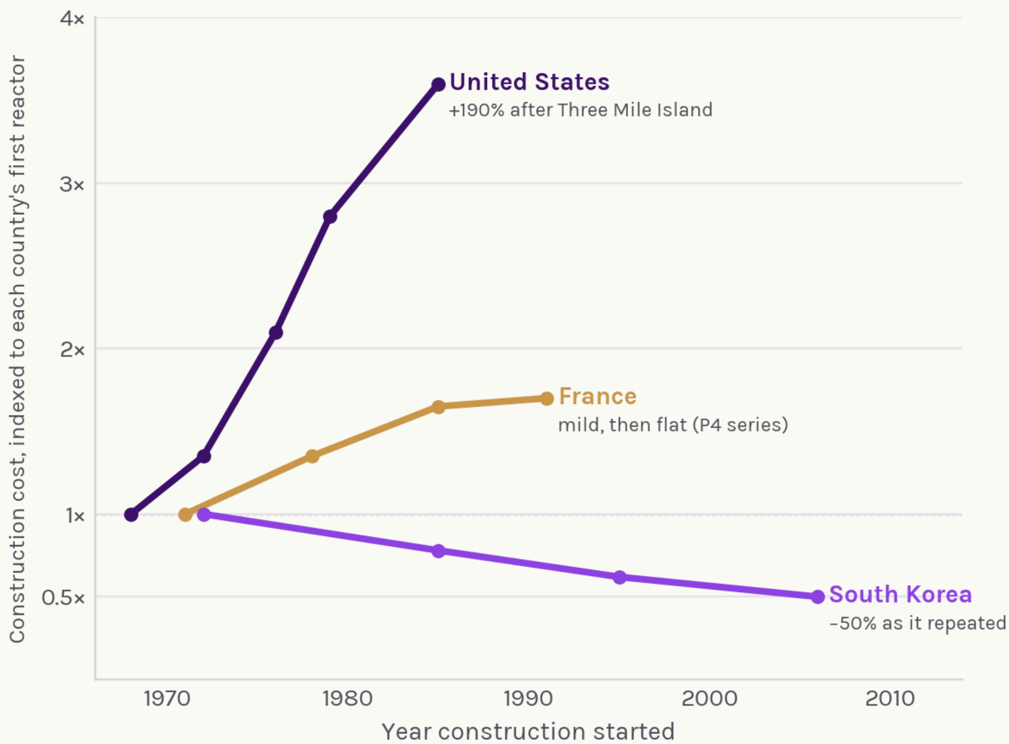

So the cause of expensive nuclear power was never purely the physics, engineering, construction, and technology. It was institutional: a particular way of building, financing, and regulating that the United States happened to adopt and Korea happened to avoid. That matters, because it means the recipe for cheap nuclear energy is already known and has been demonstrated within living memory. Standardize a design and supply chain, build it repeatedly, and hold the regulatory goalposts still. Korea did it; Japan beat a fifty-two-month schedule on its first ABWRs in the 1990s, before the fleet idled after Fukushima. The UAE, buying Korea's reactor, matched the cost at Barakah, around $4,000 per kilowatt, a quarter of Vogtle, even as the first unit slipped several years. China is doing it at scale. No modern Western build has matched any of them.

2. Why the West can't simply copy Korea

Repeating builds of large nuclear plants requires three things the West no longer has in the form Korea or China do, and each deficit pushes toward a different scale of project.

The first is a balance sheet that can absorb the cost. A Western gigawatt plant is a $20-to-60-billion, ten-year wager that, if it slips, can take its sponsor down with it. Vogtle helped bankrupt Westinghouse in 2017. Korea and China make that bet on a state balance sheet; deregulated Western utilities and merchant developers cannot, and capital markets won't underwrite a decade of single-asset construction risk. The clearest proof is the one serious American attempt to build small nuclear plants the old way. NuScale's Carbon Free Power Project, a 462-megawatt plant of six modules, collapsed in 2023 when its estimated project cost reached $9.3 billion and its municipal subscribers walked. The unit of capital the West can actually deploy is not tens of billions. It's hundreds of millions to a few billion, an order of magnitude below the gigawatt bet: a project financing, or venture rounds backed by a single customer.

The second is demand shaped like a gigawatt plant, and that's changed too. The new buyer of firm power is the hyperscaler, and behind it, the Department of Defense. What they want is power at a specific site, on a two-to-three-year clock, sized to a load that arrives in increments. A data-center campus doesn't want to wait a decade for, say, 1.1 gigawatts, and carry the merchant risk on the half it doesn't yet need. It wants tens to a few hundred megawatts, potentially even sited behind the meter, where the grid-interconnection queue (now five or more years in much of the country) doesn't apply. You can watch the shape of the demand in the deals. The binding, near-term power is mostly the *restart* of existing large plants: Microsoft's twenty-year contract for Three Mile Island Unit 1 or Amazon's colocation at Susquehanna. The genuinely *new* capacity is being contracted, mostly through non-binding frameworks rather than firm orders, from modular developers (Amazon with X-energy, Google with Kairos, Meta with Oklo and TerraPower) in 80-, 300-, and 500-megawatt blocks. The demand became modular, and the supply is following. Whether smaller can also be cheaper is the third reason, and Vogtle itself points to it.

What ballooned costs at Vogtle and Flamanville was not the reactor itself. It was the *project*: a hundred thousand bespoke tasks on a muddy site over a decade, through weather, labor turnover, financing carry, and regulations that changed mid-build. Lovering and colleagues attribute the escalation to exactly these factors, not to anything intrinsic to fission. Build the reactor in a factory, to a frozen design, and it's possible that failure mode shrinks. The factory pulls everything pullable off the critical path; what stays on site is a small, overrun-prone tenth (the seismic foundation keyed to the ground, the grid tie, the final code weld), and the hundred thousand bespoke field tasks collapse to a handful. The real prize is time. Much of what a Western reactor costs is not hardware but financing that compounds across a ten-year build; cut the build to two years and the biggest line item shrinks with it, before any learning curve is invoked. And because you build many, fast, the learning curve itself arrives in months, not the decades a gigawatt plant needs to double its output even once. This is the deep logic of going small. It's not a bet that small reactors are inherently cheaper. Instead, it's the recognition that factory repetition is the only way the West can run Asia's recipe at scale (standardization, multi-build, stable review) under its own constraints of capital and demand.

The size of the bet

A Western megaproject is a wager only a sovereign can carry. A microreactor is a venture round.

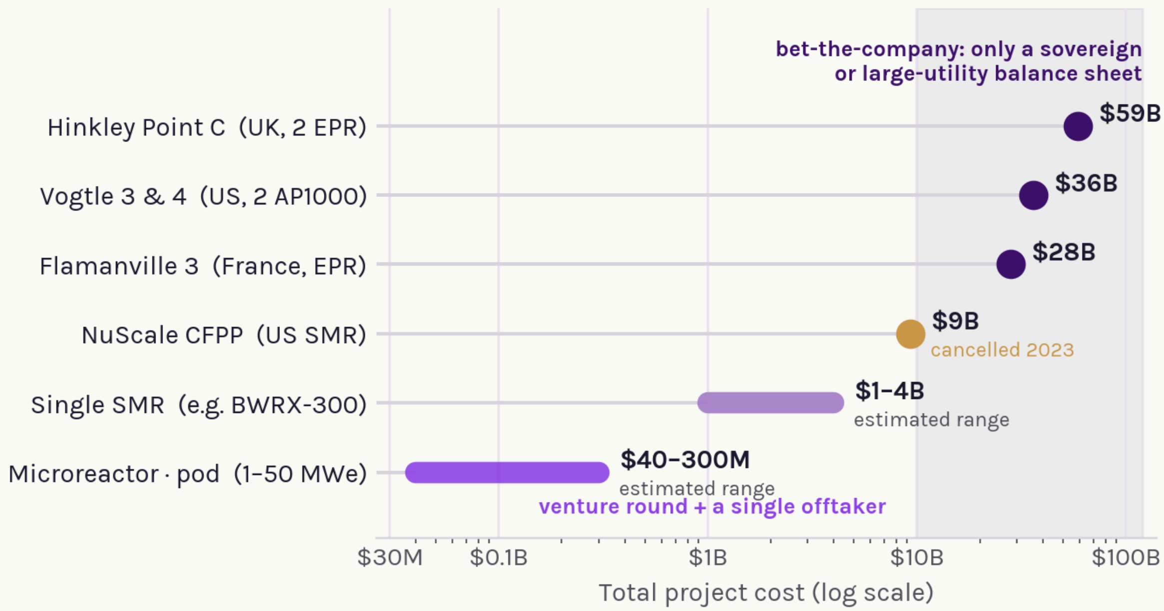

3. So how small?

For sixty years the industry's instinct was to build *bigger*, and on engineering alone the instinct is right. For a given way of building, capital cost scales roughly as the 0.6 power of output, and at a few hundred megawatts the lost economies of scale still leave a multi-unit site's operating cost per megawatt-hour roughly 20 to 25 percent above a single large plant's, a penalty on running costs, not on the all-in price of the electricity. Scale economics, in short, say build as large as you can. The trap in this line of thinking, however, is everything around the reactor. A plant big enough to capture those economies of scale is too big to leave a factory, so it becomes a site-built megaproject. And the megaproject, not the reactor, is what blows up cost. Going smaller is a deliberate trade: accept the scale penalty to buy the one advantage the West can actually execute, factory repetition.

Repetition is the advantage that pays the bill, but it's not the only thing small buys. A truckable reactor reaches where the grid won't and stands up in a couple of years, not ten. It sizes to a load that arrives in increments, sits behind the meter, and rides out a grid failure on its own. And every unit off the line teaches the next, the kind of iteration a decade-long megaproject never gets to run.

There is no single right size, though, because the thing a reactor has to beat changes with how big you build it. Idaho National Laboratory declines to name one optimum, calling the right size “more nuanced and dependent on the application.” The chart below lays the field out: cost on one axis, installation size on the other, with the competitor, in terms of levelized cost of energy (LCOE), that a plant of that respective size actually faces shaded in behind it. At the small end (a remote mine, an island, a forward operating base) the alternative is diesel at $350 to $700 a megawatt-hour, a bar so high that even a first-of-a-kind microreactor slips under it, which is why that is the beachhead. At data-center scale, the real competitor is not the grid but on-site gas at roughly $90 to $160 a megawatt-hour: reciprocating engines a hyperscaler can stand up in fifteen to twenty-four months, while frame-class gas turbines run longer, now backlogged into 2029. Only at full grid scale does a reactor face the cheap number commonly quoted, combined-cycle gas at $50 to $110, and even that sits behind a five-plus-year interconnection queue, which the behind-the-meter buyer skips, trading it for the shorter lead time of an air permit and a gas hookup. Firm renewables trace the same rough shape, expensive when small and cheap only at scale.

The competitive landscape: nuclear, gas, diesel & renewables

Approximate levelized cost of electricity vs. installation size — the bar changes with the scale of the application.

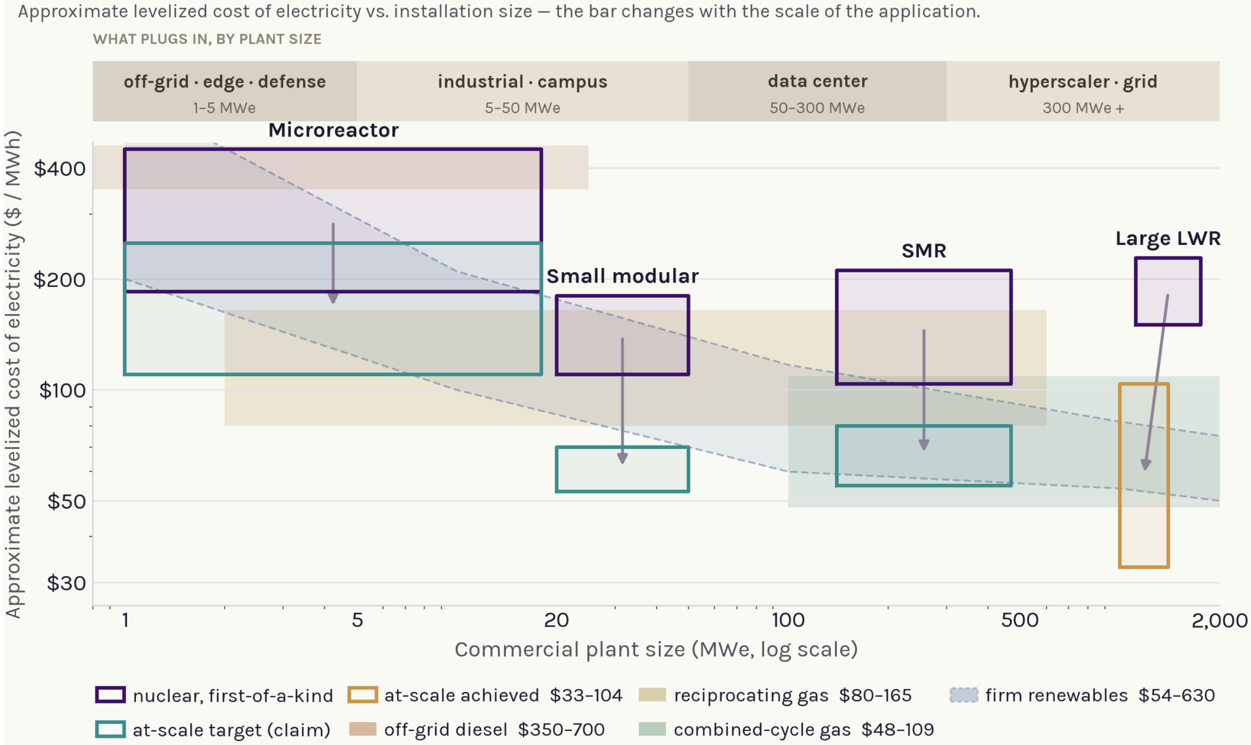

The nuclear fission field sorts into four bands by size. The labels are a moving target: GAO puts a *microreactor* under fifty megawatts and a *small modular reactor* at fifty to three hundred, but the wave of sub-megawatt, truckable machines has scrambled the neat bins, so the map above tracks the economics more than the name. At the smallest end is the microreactor, from a few hundred kilowatts to ten megawatts, road-shippable (one truckload below about 1.5 megawatts, several above). Several are designed to cluster into roughly fifty-megawatt pods;

sharing turbines across the cluster on common steam headers, with redundancy, so they can hold availability high even when one reactor is down, but it doesn't buy large-plant economics. A pod is still mostly replicated microreactors, so the bulk of the cost reduction has to come from factory volume, not from ganging more units on a site. A step up are the small modular reactors proper, roughly twenty to eighty megawatts a module, shipping a vessel and factory-built modules for assembly on site. Larger still is the SMR at the edge of road and rail, the only Western unit now under construction, at three hundred megawatts. And at the top sits a 470-megawatt machine that is an "SMR" by name but, by economics, a large reactor assembled from prefabricated parts, at the size where the modular savings thin and the burden of proof shifts to whether the bill can be financed and the second build can repeat the first.

Read against that field, only the smallest reactors, the ones racing diesel, clear their competitive bar with other energy sources today; every larger first-of-a-kind still sits above its competitor. That's a fact about first units, not a ceiling on the class: what rules a larger machine out isn't its megawatts, but a bill its sponsors can't finance and a design that never gets built twice. However, the destination the at-scale projections aim for is already occupied: Korea and China build new plants at $2,200 to $4,500 per kilowatt, roughly $50 to $70 a megawatt-hour, the only machines that ran the recipe (one frozen design, built over and over) long enough to arrive there. So the case for small does *not* rest on nuclear suddenly inheriting solar's learning curve (twenty percent cheaper per doubling); that curve has never been demonstrated for reactors. Treat the manufacturing learning rate as the upside, not the foundation. The foundation is the three structural reasons above: the West can finance a few-hundred-million-dollar reactor but not a $20-to-60-billion one; its demand arrives in increments, not gigawatt blocks; and the project, not the physics, is what blew up Vogtle. The learning curve is the bet on top, and, as we'll see, the cheapest bet of its kind ever offered.

4. Could it actually be cheap

Suppose repetition gets the West to Barakah's ~$4,000 per kilowatt. Cheap to build still doesn't settle it. Duan, Petroski, Wood, and Caldeira (*Nature Energy*, 2022) found that even at $4,000 (below today's cost), wind and solar stay cheaper under *moderate* decarbonization; nuclear pulls ahead only past about 80 percent emissions reduction, where cheap storage can no longer cover the gaps. Necessary, not sufficient.

That reads like a refutation against nuclear. Instead, it's the hinge back to the first piece: AI is manufacturing exactly the condition under which nuclear wins, a requirement for firm, around-the-clock power that storage can't yet fill, on a grid decarbonizing under a load that only grows as society demands more compute. However, the near-term bar is gentler than Duan's deeply decarbonized grid. A reactor behind a data center isn't racing an eighty-percent-clean grid; it's racing the on-site gas from the chart above, $90 to $160 a megawatt-hour, which an at-scale reactor near $60 beats. Own the gap, though: $60 is the destination, not today's price. A first-of-a-kind at $10,000 per kilowatt pencils in nearer $150 a megawatt-hour, Vogtle territory, which is exactly why the walk-down, not the first unit, is the whole game. The case for cheap nuclear was never that it beats solar at noon. It's that, built like a product, this is some of the cheapest firm power you can put almost anywhere, running around the clock and refueled only occasionally.

One caveat belongs here. Levelized cost is how you compare machines, not how power gets bought. The buyer signing the check cares where the kilowatt-hour lands, when it can be energized, and how firm it is (the agencies that publish levelized costs also publish value-adjusted versions for exactly this reason), and the project itself turns on an air permit and on financing, because an investment-grade offtake compresses the cost of capital, and capital is roughly half the cost of firm power. The chart above says which reactors are worth comparing. Buyers and bankers decide which get built. NuScale failed both tests at once: its price climbed from $55 to $89 a megawatt-hour, which the cost metric caught, and the project was cancelled anyway, because no one would finance the schedule.

5. The race to criticality

The fastest way to see which approaches are real is to watch them try to start. The Department of Energy's Reactor Pilot Program set a public, dated bar: at least three privately built reactors critical (a chain reaction becoming self-sustaining) by July 4, 2026. Four cleared it.

Part of what makes that pace possible is a quieter accelerant. Some of the smaller teams are running AI through the engineering itself, from design iteration to cost models to licensing paperwork, and they are adopting it faster than the established players. The tools alone will not decide the outcome, though. These are compound problems (hard engineering, manufacturing, regulation, and finance, all at once) that will take everything available, and the likely winners are the teams that pair disciplined engineering management with the best tools rather than betting on either alone. Either way, the same technology creating the demand for these reactors is compressing the time it takes to engineer them.

Criticality, however, is not an electron. The article that goes critical can be as little as fuel, a moderator, and control drums at zero power; it proves the physics, but it doesn't light a bulb, run through a winter, or power a data center. The earliest a privately developed Western reactor is targeted to generate any electricity is 2027, and that's a demonstration unit at a national lab, barred from selling to the grid, not a commercial plant. First customer electricity is likely a 2028 event; commercial operation at scale belongs to the 2030s. Fuel is the other gate: most advanced designs need high-assay low-enriched uranium, the fuel the West barely makes yet, so the light-water SMRs that run on ordinary sub-five-percent uranium have the clearest path. And even once enrichment is solved, exotic fuel forms like TRISO still have to be fabricated and qualified (the first commercial TRISO line was licensed only in early 2026), so fuel is a question of timing, around 2029 to 2031, not physics.

6. Why now — and why the door is opening

The framework that licensed the entire American fleet, Part 50, split approval into a construction permit and then a separate operating license, a two-step gauntlet; even the streamlined one-step license that replaced it, Part 52, took the Nuclear Regulatory Commission about four years to grant Vogtle, inside a fifteen-year march from application to power. The NRC's new technology-inclusive framework, Part 53, took effect in April 2026 and licenses a reactor on its actual consequence rather than on the assumption that every reactor is a gigawatt pressurized-water plant. Even the permits moving under the older rules have come in far faster: Kairos's Hermes took roughly two years. And in May 2025, Executive Order 14300 ordered the NRC to set hard deadlines where it had effectively had none: eighteen months for a final decision on a new reactor, twelve for an existing one's continued operation. Running alongside, the Department of Energy's pilot program is putting reactors on the ground under federal authorization, outside the NRC's licensing queue entirely.

The newest designs are safer by *physics*, not by adding backup systems. Shutting down the fission was always the relatively easy part; the harder part is the decay heat that keeps coming for days after: six to seven percent of full power at the instant of shutdown (the heat that melted Fukushima). The newest designs reject it passively, through conduction and convection, because a small core has the surface-area-to-volume ratio to lose its heat without a pump. (Even large Gen III+ plants like the AP1000 now ride out seventy-two hours with no power; shrink the core and that window stretches toward indefinite.) A reactor that loses all electricity and all operators shuts down and stays cool on its own, and "walk-away safe" is what lowers the consequence envelope enough to make consequence-based licensing defensible. The three accidents that define public memory were failures of old designs and active systems: Three Mile Island released little radiation and killed no one; Chernobyl was a Soviet design with a positive void coefficient and no real containment, both engineered out of every Western reactor; Fukushima, a 1960s plant whose cooling pumps lost power to a tsunami, caused one radiation death, and the evacuation did far more damage than the radiation. On a deaths-per-terawatt-hour basis, one that counts the full Chernobyl and Fukushima tolls, evacuation included, nuclear sits beside wind and solar as the safest power humans have built.

Among the safest power we have

Deaths per unit of electricity, accidents and air pollution included. Nuclear in purple.

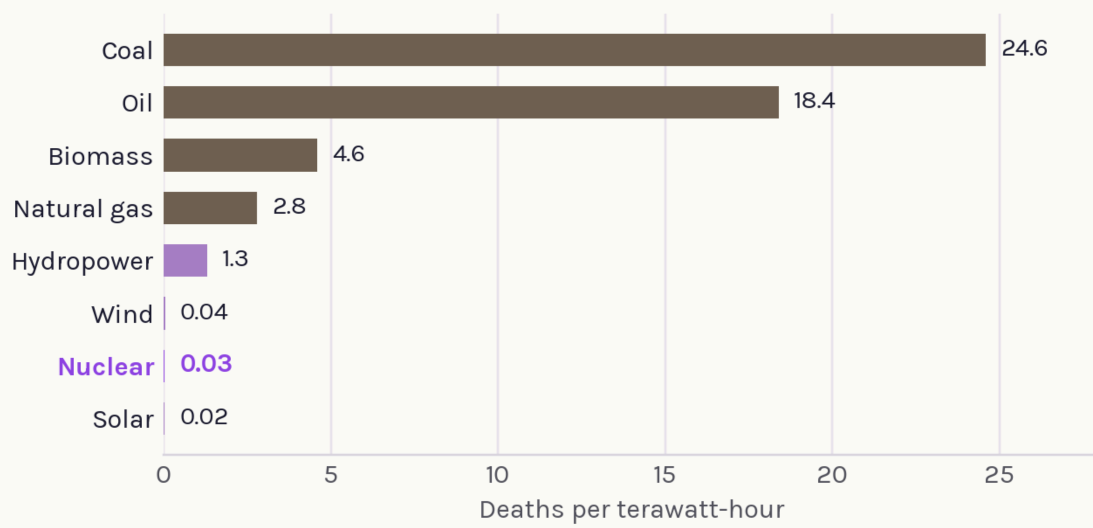

7. The bet worth making

So read the $36 billion for what it is: the price of a megaproject, not a verdict on the reactor. The firm-power market won't go to the most elegant design, since most of them work, but to whoever stops building megaprojects and starts building products. That is the only form in which the West can finance repetition, serve the demand it actually has, and escape the failure mode that made nuclear expensive in the first place.

None of it is guaranteed. It needs the West to relearn serial production at a size its capital markets can fund, the fuel buildout to land close to schedule, and the federal subsidies that make the first at-scale numbers pencil to outlast a construction clock that runs into the 2030s, all while the same hyperscalers underwriting it can already turn to behind-the-meter gas today. So watch the one signal that matters more than any groundbreaking or press release: cost falling across the first several reactors off a production line, not the promise of it on the first. Until that day this is a conviction. After it, an industry.

If you're working on any of this, or on anything one architectural layer removed from the reactor itself, I'd like to hear from you. roger@primemoverslab.com.

Sources & method. Cost figures are overnight or total-project cost as publicly reported and are not inflation- or scope-harmonized; multiples compare final or latest cost to the budget at sanction. Vogtle's ~$16,000 per kilowatt is derived from the ~$36 billion total over the two-unit, ~2.2-gigawatt build; treat per-kilowatt comparisons to Barakah and the Chinese fleet as order-of-magnitude. The $/MWh figures are PML estimates at a 90% capacity factor and ordinary financing, and are illustrative; at a bankable ~75% capacity factor and with the ITC/PTC stripped, the at-scale figures rise roughly $25–40 a megawatt-hour. The financing-wall figures for a single SMR and a microreactor are illustrative order-of-magnitude ranges. Project costs: Georgia Power / GAO (Vogtle); EDF and the Cour des comptes (Flamanville, €23.7B / ~$28B, 2025); TVO (Olkiluoto); UK National Audit Office and New Civil Engineer 2026 (Hinkley, ~£46B / ~$59B); Barakah and the Chinese fleet from the World Nuclear Association and NucNet; NuScale Carbon Free Power Project ($9.3B, 462 MW, cancelled 2023) from UAMPS / NuScale. Cost trends and the "no inherent escalation" finding: Lovering, Yip & Nordhaus, Energy Policy 91 (2016); the "negative learning" reading it corrects is Grubler, Energy Policy 38 (2010). Size–cost trade: Carelli et al., Progress in Nuclear Energy (2010, scale exponent ~0.6); Mignacca & Locatelli, RSER (2020); Hanna, Abou-Jaoude et al., Energies 17(10):2390 (2024, INL — application-dependent optimum); Abdulla, Azevedo & Morgan, PNAS 110(24):9686 (2013); Carelli et al. 2010 (multi-SMR site O&M roughly 22–24% above a single large reactor). Demand & lead times: company announcements 2025–26 (Microsoft / Constellation; Amazon / Talen; Amazon / X-energy; Google / Kairos Power; Meta / Oklo; Meta / TerraPower (up to 8 Natrium units, January 2026); Meta / Vistra), capacities as announced, most new-build agreements frameworks or letters of intent rather than firm PPAs; gas lead times consistent with the first piece in this series — reciprocating engines ~12–24 months, frame-class turbines backlogged into 2028–2029 (GE Vernova ~100 GW backlog, Siemens Energy sold out through 2028; Utility Dive / Jefferies, Feb 2026), aeroderivatives 18–36 months. Cost-competitiveness: Sepulveda et al., Joule (2018); Duan, Petroski, Wood & Caldeira, Nature Energy 7(3):260–269 (2022), DOI 10.1038/s41560-022-00979-x — even at $4,000/kW, wind+solar cheaper under *moderate decarbonization, nuclear ahead past ~80% emissions reduction. Competitive-landscape chart: benchmark LCOE bands from Lazard LCOE+ (June 2025), with NREL ATB, EIA and SemiAnalysis for off-grid diesel, reciprocating and combined-cycle gas, behind-the-meter lead times, and firm renewables; reactor boxes span PML per-company estimates by size class — first-of-a-kind from independent estimates, at-scale from developer targets (for microreactors, production-cost targets of roughly high-teens to mid-20s ¢/kWh — a cost, not a sellable PPA rate) except the large-reactor fleet box, which is achieved operating data (Korea and China builds, the UAE's Barakah, after OECD Projected Costs 2020 and NEI/EIA). SMR first-of-a-kind cost ~$10,000/kW: International Energy Agency, "The Path to a New Era for Nuclear Energy." Reactor Pilot Program outcome: EO 14301 set a goal of at least three test reactors critical by July 4, 2026; four achieved it — Antares Nuclear Mark-0 first, June 4; Valar Atomics, June 18; Deployable Energy; Aalo Atomics, July 4 (Department of Energy; World Nuclear News; POWER; American Nuclear Society). Mark-0 is a zero-power test reactor (no power rating); the commercial unit is rated 100 kWe–1 MWe. First demonstration electricity targeted for 2027 (Antares Mark-1 at INL; Oklo Aurora-INL, barred from grid sales); first customer electricity a 2028 milestone by developer roadmaps. Fuel: HALEU is 5–20% enriched vs. sub-5% LEU; TRISO fabrication licensed Feb 13, 2026 (X-energy / TRISO-X SNM-7007, first new NRC fuel-fab license in ~50 years) but the fuel is not yet NRC-qualified for commercial use; enrichment/conversion buildout (DOE $2.7B task orders, Jan 2026) and the Russian-import ban (absolute Jan 1, 2028) push the binding fuel timeline to ~2029–2031. Licensing: existing fleet under Part 50's two-step process; Part 52 combined license took the NRC ~4 years for Vogtle inside a ~15-year cycle; Part 53 effective April 2026; recent advanced-reactor construction permits far faster (Kairos Hermes ~16–25 months); Executive Order 14300 (May 23, 2025) directs the NRC to issue final decisions within 18 months for new reactors and 12 months for existing-reactor continued operation; DOE Reactor Pilot Program under Executive Order 14301. Japan ABWR on-schedule datum (Kashiwazaki-Kariwa 6 & 7, on/ahead of a 52-month target, 1996–97): OSTI; NEI. AI-accelerated engineering: author conversations with reactor developers, established and startup, 2025–26. Naarea liquidation: French commercial-court registry and L'Usine Nouvelle. NuScale CFPP subscription price rose $55 to $89/MWh, 2021 to January 2023, including the federal PTC (UAMPS / NuScale). Value-adjusted cost metrics: EIA levelized avoided cost (LACE); IEA value-adjusted LCOE (VALCOE). Financing share of Western build cost: MIT, The Future of Nuclear Energy in a Carbon-Constrained World (2018); OECD/NEA, Unlocking Reductions in the Construction Costs of Nuclear (2020).*