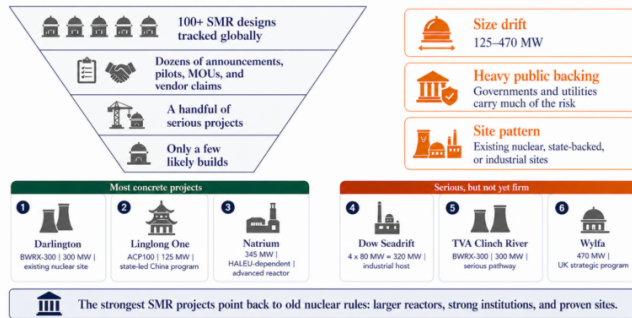


SMR Reality Check: Few Builds, Bigger Reactors, Proven Sites

The market is producing a handful of larger, heavily backed projects.



Infographic of SMR reality by author with ChatGPT

Nuclear Scaling Requires Discipline. SMRs Deliver Fragmentation

17 hours [Michael Barnard](https://cleantechnica.com/2026/04/28/nuclear-scaling-requires-discipline-smrs-deliver-fragmentation/) 11 Comments <https://cleantechnica.com/2026/04/28/nuclear-scaling-requires-discipline-smrs-deliver-fragmentation/>

When I wrote in 2021 that small modular reactors were mostly bad policy ([peer reviewed version](#), [CleanTechnica version](#)), the argument was not that nuclear fission could not produce useful low-carbon electricity. It was already doing so every day. The United States had about 98 GW of operating nuclear capacity, and the global fleet was a major source of firm generation. The question was whether the SMR policy proposition matched the conditions under which nuclear power had scaled in the past. It did not then. The evidence since then has made the problem clearer.

The original SMR case rested on a simple promise. Make reactors smaller, build more of them in factories, reduce capital at risk, shorten construction schedules, serve more sites, and avoid the large-project failures that had damaged recent nuclear construction in liberalized electricity markets. It was an appealing story because it pointed at real nuclear problems. Large reactors are expensive to finance. They take a long time to build. A single failure can consume a utility's balance sheet and a government's political patience. A smaller unit sounds easier to manage.

But the promise depended on a condition that was often treated as background noise. SMRs only make economic sense if the sector converges on a few designs and builds them many times. Factory manufacturing does not create a learning curve because the word factory appears in a presentation. Learning curves come from repeated production of the same or similar products, with stable tooling, stable suppliers, stable inspections, stable quality assurance, stable training, and steady demand. Solar panels, batteries, and wind turbines became cheaper because the world made huge numbers of related products in shorter production cycles. Nuclear reactors are different. Each design carries a safety case, a fuel qualification pathway, licensing work, site work, security, emergency planning, operator training, waste arrangements, and decades of liability.

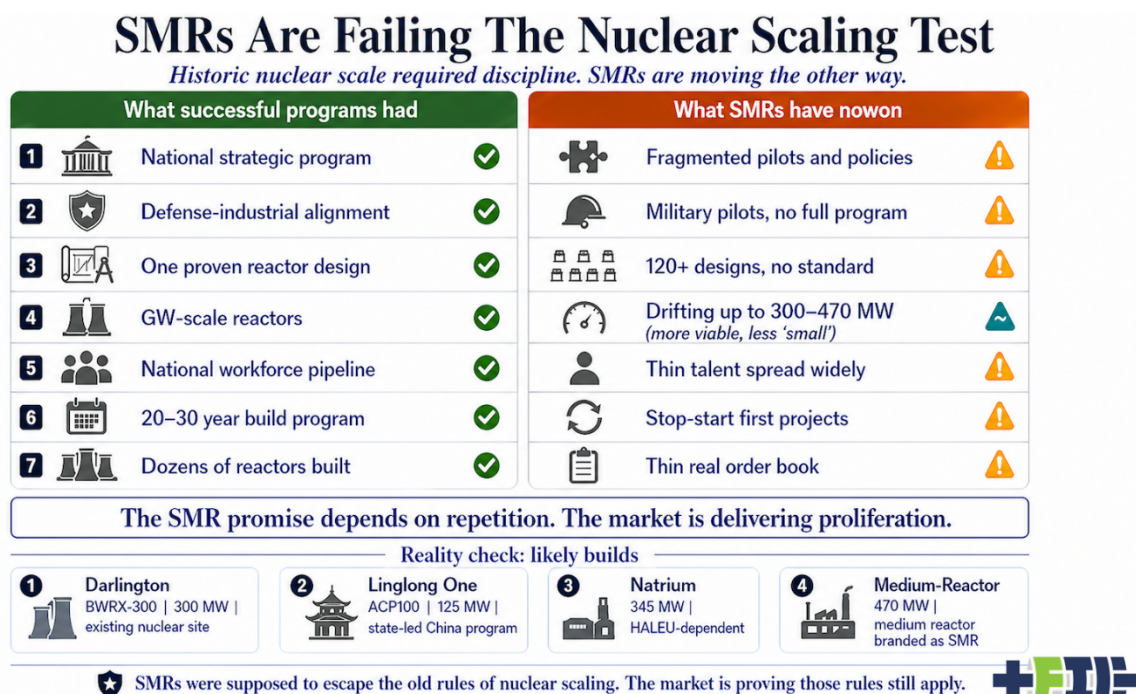
That was the central weakness in the SMR story in 2021. In that earlier assessment, I counted 57 SMR designs and concepts across 18 broad types, and none could be considered dominant. That was already far too fragmented for a credible manufacturing-learning argument. Since then, the OECD Nuclear Energy Agency's SMR dashboard has tracked more than 120 SMR technologies worldwide, with roughly 70 to 80 included in recent dashboard editions after filtering out some

paused, inactive, unfunded, or non-participating designs. The sector has not moved from many concepts to a few winners. It has become more crowded.

This matters because nuclear design proliferation is not cheap experimentation. In software, a hundred teams can try different approaches, fail fast, and leave lessons behind. In nuclear, each credible design requires scarce engineering, regulatory, fuel-cycle, owner, and supply-chain attention. A light-water SMR, a high-temperature gas reactor, a sodium fast reactor, a molten-salt reactor, and a microreactor are not minor variations around a shared product platform. They create different materials questions, fuel requirements, operating temperatures, inspection regimes, safety cases, and licensing pathways.

The [EIA's April 2026 Today in Energy article](#) is useful because it lays out that diversity. It groups U.S.-relevant SMRs and microreactors into light-water reactors, high-temperature gas reactors, molten-salt reactors, sodium-cooled reactors, and other designs. It identifies applications such as AI loads, data centers, industrial sites, remote areas, microgrids, and military or federal facilities. It points to DOE programs, pilot pathways, and fuel-chain efforts. As a map of activity, it has value. As a test of whether the SMR proposition is becoming a real deployment class, it is much weaker.

The EIA article does not ask the questions that matter for scaling. It does not ask whether the order book is large enough to support factory learning. It does not ask whether design proliferation undermines standardization. It does not ask whether the credible projects are really small, or whether they are drifting back toward conventional power-station scale. It does not ask whether remote sites, mines, and islands are large enough markets to sustain a reactor manufacturing industry. It does not ask whether HALEU will be available at scale on the timelines implied by advanced reactor plans. It describes activity and optionality. It does not demonstrate convergence.



Infographic testing SMRs against nuclear scaling success criteria by author with ChatGPT

The historical conditions for nuclear scaling are not mysterious. Nuclear built at scale where it

were standardized or semi-standardized, where large reactors spread fixed costs over a lot of output, where experienced nuclear owner-operators existed, where training and safety culture were centralized, and where governments sustained programs for decades. France, South Korea, and China did not scale nuclear power by letting dozens of small reactor startups compete for scattered boutique sites. They scaled, to the extent they did, through alignment among state policy, utilities, vendors, regulators, finance, and workforce.

SMRs were sold as a way around these conditions. The actual market is rediscovering them. The projects that look most likely to be built are tied to existing nuclear sites, state-backed strategic sites, experienced utilities, military or laboratory settings, or large industrial anchors with public support. That does not mean they are worthless. It means they are not validating the broad SMR pitch. They are validating the old lesson that nuclear needs strong institutions.

The most credible projects are also getting bigger. Ontario's Darlington project is the clearest Western example. Ontario Power Generation has a license to construct one GE Hitachi BWRX-300 at Darlington, with four units planned. Each unit is about 300 MW. This is a serious project, but it is not a small reactor scattered into a new class of sites. It is a 300 MW boiling water reactor at an existing nuclear site, backed by an experienced provincial nuclear operator with grid interconnection, cooling access, security culture, political support, and a long-term system need. If it succeeds, it will matter. But it will not prove that SMRs can escape nuclear's institutional requirements.

China's Linglong One, the ACP100 at Changjiang in Hainan, is another real project. At about 125 MW, it is closer to the traditional idea of a small reactor, and it has moved through construction and testing milestones. But it exists inside China's state-led nuclear program. China can choose, license, finance, build, and integrate nuclear projects in ways that liberalized markets struggle to copy. That makes Linglong One important, but it does not make it proof that a global commercial SMR market has arrived.

TerraPower's Natrium project in Kemmerer, Wyoming, is serious as well, with a construction permit issued by the U.S. Nuclear Regulatory Commission and non-nuclear site work underway. But Natrium is 345 MW, with storage-boosted output advertised around 500 MW. It sits above the old 300 MW SMR threshold and depends on sodium cooling, HALEU fuel, major public support, and a coal-site transition narrative. It may become a useful advanced reactor demonstration. It is not evidence that small, repeatable, low-risk nuclear products are ready for broad deployment.

Rolls-Royce makes the size drift even more obvious. Its reactor is about 470 MW. Three units at Wylfa would total about 1.4 GW, which is a large power station by any normal electricity-system measure. The unit is small only compared with the largest conventional reactors. It may fit the United Kingdom's industrial strategy if the government commits to a fleet. But at 470 MW, the project is better understood as a medium reactor with modular construction ambitions than as the small product implied by early SMR rhetoric.

Holtec's design history points the same way. The SMR-160 became the SMR-300. NuScale's module moved from 50 MW toward 77 MW, and the commercial plant concept became a multi-module station approaching conventional plant scale. X-energy's Xe-100 is about 80 MW as a

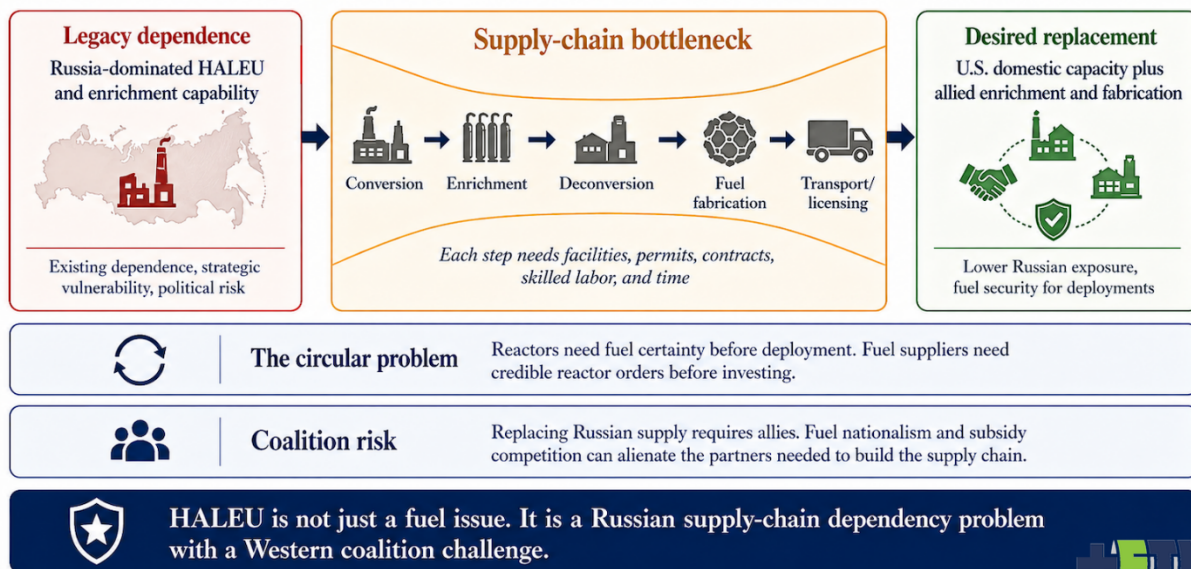
pattern is clear. The more serious the customer discussion becomes, the more the sector tries to put several hundred MW behind a single site, operating organization, licensing file, security plan, and grid connection.

That is not an accident. Nuclear has large fixed costs that do not shrink in proportion to reactor size. A 50 MW reactor does not need one-twentieth of the licensing effort, one-twentieth of the security analysis, one-twentieth of the operator training, one-twentieth of the emergency planning, one-twentieth of the quality assurance, or one-twentieth of the waste arrangements of a 1,000 MW reactor. Some hardware costs scale down. Many institutional costs do not. Smaller reactors start with a scale penalty. Factory repetition is supposed to overcome it. But repetition requires a narrow set of designs and a large order book. The current market offers neither.

After years of SMR hype, the likely-build list remains short: Darlington, Linglong One, Natrium in Wyoming, TVA's Clinch River, Dow's Seadrift project, Holtec's proposed Palisades units, Rolls-Royce at Wylfa, and Russian RITM-based Arctic or floating projects. That is not nothing, but it is not a broad commercial market. It is a small order book of state-backed, utility-backed, or strategic projects, often tied to existing nuclear or heavy-industrial sites, often larger than the original SMR story implied, and often dependent on public risk absorption. By contrast, the press-release order book is filled with memoranda of understanding, technology selections, data-center announcements, export discussions, remote-site narratives, and vendor road maps. Those are not reactors. Nuclear projects have a long valley between interest and electrons.

HALEU: The Fuel Bottleneck Behind Advanced Reactors

Replacing Russian supply requires a Western fuel chain that is not yet commercial at scale



Infographic of HALEU advanced reactor fuel bottleneck by author and ChatGPT

HALEU sits near the center of the problem, not at the edge of it. Several advanced reactor designs require higher-assay low-enriched uranium, enriched above the 3% to 5% U-235 common in today's light-water reactor fuel but below 20%. HALEU can support smaller cores, longer operating cycles, higher burnup, and reactor designs that standard low-enriched uranium cannot support. That is why developers want it. It is also why it is a bottleneck.

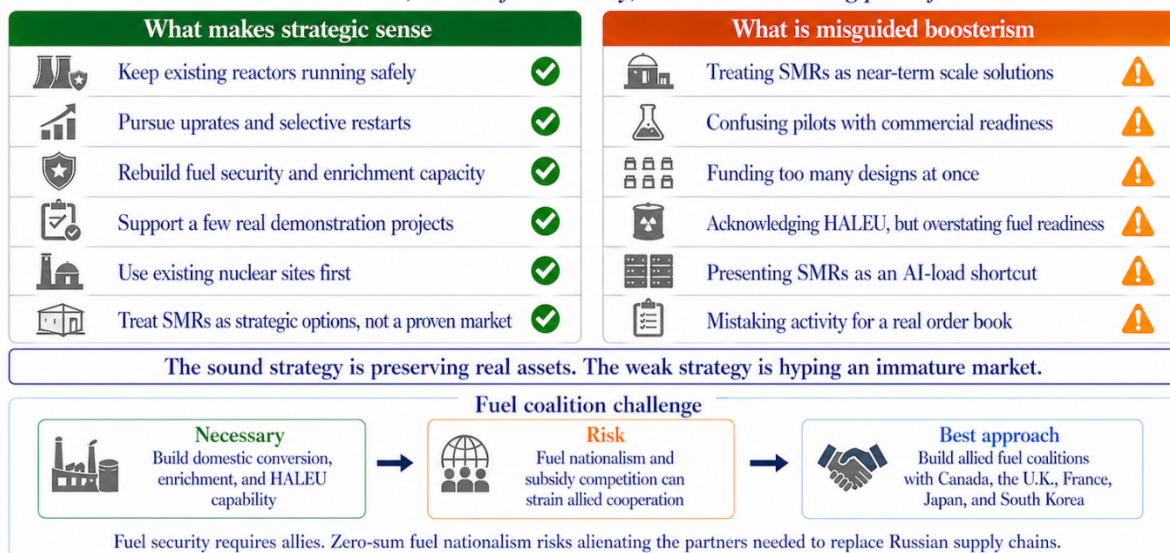
The United States does not yet have a mature, large, domestic HALEU supply chain. Russia has been the major commercial source, which is now a strategic and political problem. Rebuilding a domestic chain requires conversion, enrichment, deconversion, fuel fabrication, transport packages, licensing, inspections, safeguards, workforce, and customer commitments. Each link needs facilities, capital, permits, contracts, and time. This is not a paperwork problem. It is an industrial-base problem.

There is a circular dependency at the heart of it. Reactor developers need HALEU to make credible deployment commitments. Fuel suppliers need credible reactor demand to justify investment. Customers need confidence that both reactor and fuel will be available. Regulators need data on fuel behavior and safety. Government can break pieces of the loop by funding fuel production and demonstration quantities, but that confirms that the strategy is government-led. It does not show that advanced SMRs are market-ready.

HALEU also makes design proliferation more damaging. A narrow reactor program using a common fuel form creates a clearer demand signal. A market with many designs, fuel forms, enrichments, geometries, claddings, coolants, and operating conditions creates a harder investment problem. Fuel suppliers are not being asked to serve one standardized fleet. They are being asked to prepare for a moving set of possible reactor futures. If HALEU is a gating condition for deployment, then public policy should be narrowing the field, not celebrating breadth.

U.S. Nuclear Strategy: Sound Policy vs SMR Boosterism

Preserve real assets, rebuild fuel security, and avoid mistaking pilots for scale.



Disciplined, fuel-secure, and allied beats aspirational SMR strategy.



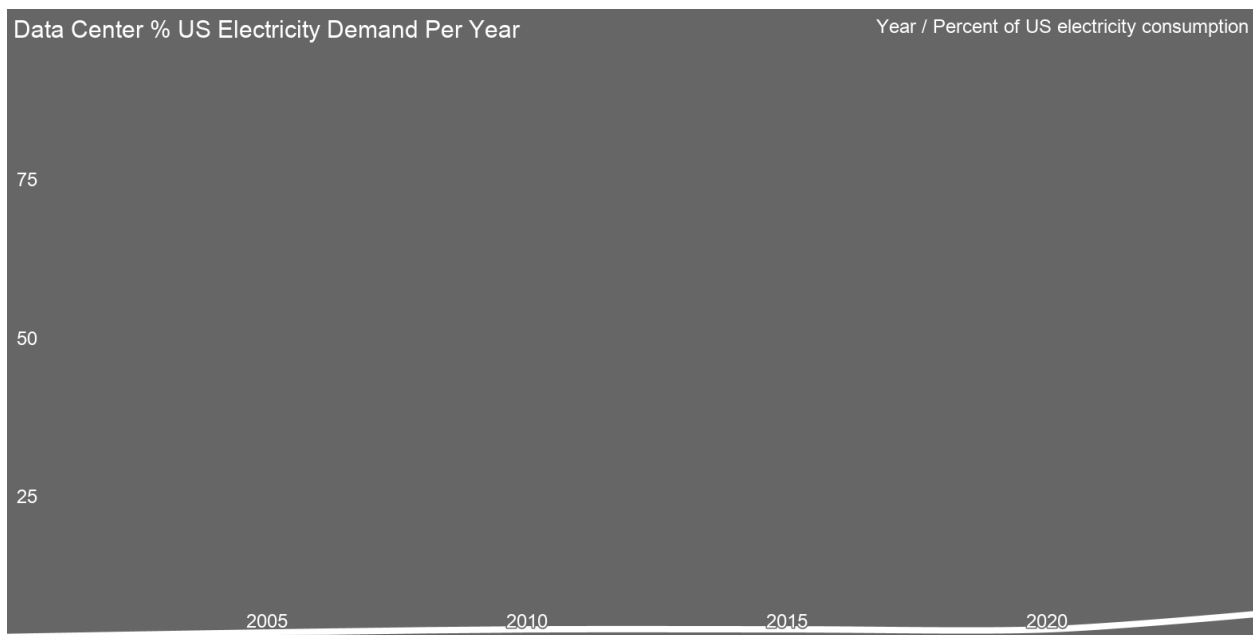
Infographic assessing United States nuclear energy strategy by author with ChatGPT

This is where U.S. energy policy becomes confused. The United States has a rational nuclear policy layer and a speculative nuclear policy layer. The rational layer is preserving safe existing reactors, extending licenses where appropriate, uprating existing units, restarting recently retired units where the equipment and economics support it, and strengthening the workforce and fuel system. Existing plants have grid connections, trained operators, known safety records, community relationships, cooling systems, and regulatory histories. Keeping a safe reactor

operating can avoid large volumes of fossil generation with much less uncertainty than a first-of-a-kind new build.

The speculative layer is treating a fragmented SMR sector as if it were already a deployable answer to new load growth. DOE's UPRISE initiative, which emphasizes uprates, restarts, license extensions, and improvements to existing reactors, belongs largely in the practical bucket. A \$900 million Gen III+ SMR funding opportunity belongs in the option-value and industrial-policy bucket. It may help one or two designs move forward. It may produce learning. But it is not proof that the commercial case exists.

AI has become the new accelerant for this policy story. Data centers want large amounts of firm power, often on fast schedules. U.S. policymakers are concerned about electricity demand growth from AI, data centers, and advanced manufacturing. Nuclear advocates see an opening. The problem is timing. Data centers are being planned and built on two-year to five-year horizons. First-of-a-kind nuclear projects move through design completion, licensing, site work, supply-chain development, fuel procurement, construction, testing, and commissioning on longer timelines. Existing nuclear plants can serve some corporate procurement needs. Restarts and uprates may help in some places. SMRs are not close enough to be the main answer to near-term AI load.



Data Center % US Electricity Demand Per Year by author

Data centers are a shaky foundation for SMR strategy in any event because the AI electricity panic has already started to look familiar. As I argued in a [January 2025 CleanTechnica piece](#), every wave of digital growth has produced claims that data centers were about to overwhelm the grid, from the dot-com boom to cloud computing, streaming, remote work, blockchain, and now AI. The pattern has been repeated concern, then hardware, software, architecture, and market optimization. U.S. data centers were about 1.5% of electricity consumption in the 2006 EPA report and only about 1.8% in 2014, despite the internet becoming central to daily life. Even with AI, the article noted data centers at about 4.4% of U.S. electricity demand in 2022, material but not world-ending.

The more important point is that AI efficiency is already improving fast: Nvidia's Blackwell architecture claimed up to 25-fold better energy efficiency for inference than the prior generation, while DeepSeek showed how software optimization could deliver comparable model performance with much lower compute costs. Jevons Paradox means cheaper AI may still increase total use, so the demand spike is not imaginary. But it is volatile, innovation-sensitive, and easy to exaggerate. Using AI load growth to justify first-of-a-kind SMRs confuses a moving digital demand story with a proven reactor market. Data centers need power on short commercial timelines. SMRs still need design completion, licensing, fuel, construction, commissioning, and repeat orders.

That is the core policy failure. U.S. SMR policy is confusing aspiration, option value, and industrial strategy with deployment readiness. Policymakers want SMRs to support AI growth, military resilience, export competition, coal-site redevelopment, industrial heat, fuel-cycle rebuilding, and decarbonization before the sector has demonstrated cost, schedule, fuel readiness, repeat construction, or customer depth. That is misguided boosterism. It takes a category that should be treated as a narrow, risky, publicly supported technology option and presents it as if it were a near-term pillar of energy strategy.

Microreactors and remote-site claims should be separated from utility-scale SMRs. Military bases, national laboratories, and research campuses are credible early niches because they have strategic reasons to accept higher cost, unusual risk, and federal procurement structures. Project Pele at Idaho National Laboratory, a 1 MW to 5 MW transportable reactor demonstration for the Department of Defense, fits that category. It is strategic procurement. It is not evidence of normal commercial electricity competitiveness.

Remote communities, mines, and islands are weaker as broad markets. They have real energy problems, including high diesel costs, reliability challenges, fuel logistics, and limited grid access. But the alternatives are improving and being built now. Mines in Western Australia have deployed hybrid systems with solar, wind, batteries, controls, demand management, and gas or diesel backup. Gold Fields' Agnew project has delivered roughly 50% to 60% renewable energy over the long term. Lioneville's Kathleen Valley project targets more than 60% renewable power from startup. Those systems are modular, financeable, serviceable by normal industrial contractors, and expandable in pieces. They do not require nuclear licensing, nuclear operators, HALEU supply, nuclear waste arrangements, or a nuclear security regime.

The same logic applies to islands and remote communities. Solar, wind where resources are good, batteries, thermal storage, demand response, efficiency, heat pumps, and retained backup can reduce fuel imports and improve resilience without importing the full institutional weight of a nuclear facility. A microreactor may make sense for a sovereign military site, a national laboratory, or a nuclear-capable jurisdiction with a strategic reason to pay for it. That is different from a scalable business model. When an energy technology retreats to remote sites as a leading commercial story, it is often no longer arguing that it is broadly competitive. It is arguing that unusual constraints may hide its disadvantages.

A rational policy would stop treating optionality as progress. If governments believe SMRs are strategically necessary, then they should fund discipline. Pick one or two designs for fleet

reporting. Separate first-of-a-kind cost from claimed nth-of-a-kind cost. Tie public support to standardization, real orders, fuel readiness, and repeat construction. Do not count MOUs as demand. Do not pretend that every data-center press release is a reactor order.

Licensing reform can help, but it is not a substitute for a market. The ADVANCE Act and related U.S. efforts to make NRC processes more timely and predictable are reasonable in principle. Regulators should be efficient while maintaining safety and security. But if dozens of designs seek attention, faster licensing does not solve the deeper problem. The bottleneck moves to design maturity, fuel, supply chain, owner capability, financing, construction execution, and public acceptance.

The policy mistake is not supporting any SMR development. Governments often buy option value, and there can be reasons to maintain nuclear engineering capacity, preserve strategic fuel-cycle skills, support a few demonstrations, and keep an export option alive. The mistake is presenting a fragmented, fuel-constrained, thinly ordered technology class as if it were a central answer to near-term electricity demand, AI growth, or industrial decarbonization. That is boosterism, not rational energy policy.

The update to the 2021 conclusion is straightforward. The success conditions have not been met. The sector has not consolidated. The credible projects are getting larger. The real builds are mostly attached to existing nuclear sites, state-backed programs, or strategic industrial contexts. HALEU remains a hard constraint. Remote-site narratives remain niche claims. Small, modular, advanced, factory-built, flexible, and deployable are claims that have to survive contact with licensing, fuel, siting, security, staffing, waste, construction, financing, and repeat orders. Some reactors will likely be built. Some may be useful. But the evidence does not support treating SMRs as a broad, near-term, commercially validated solution. It supports the older and less exciting conclusion that nuclear scale requires focus, standardization, strong institutions, mature fuel supply, and a long program. The SMR sector is still moving in the opposite direction.