Small Modular Reactors

The factory-built bet to make nuclear cheap, fast, and flexible

Explainer · June 6, 2026

ABSTRACT

Russia's floating Akademik Lomonosov has been feeding power to a remote Arctic town since May 2020, and China's pebble-bed HTR-PM entered commercial service in 2023, but the West has yet to switch on a single small modular reactor (SMR). These are nuclear fission reactors of up to roughly 300 megawatts of electrical output, about a third the size of a conventional plant, whose central wager is manufacturing rather than physics: build major components in a factory, ship them as modules, and assemble on site to compress schedules and control quality. Designs and regulatory approvals are now maturing into the first Western construction projects. Yet the decisive question is economic, not technical: whether serial production can drive costs down from punishing first-unit estimates to competitive levels. Implications run from data-center power deals to a global pipeline of 127-plus designs against only about seven reactors built or building.

Keywords: small modular reactors; nuclear fission; modularity; passive safety; HALEU; NuScale; levelized cost of electricity; data centers

1. Introduction

The signal moment is regulatory and physical at once. In May 2025, U.S. regulators approved NuScale's uprated 77-megawatt module, building on the company's earlier first-ever U.S. design approval for a small modular reactor. Meanwhile, two flagship Western projects have moved from paper toward steel: TerraPower's 345-megawatt Natrium plant in Wyoming and GE Hitachi's BWRX-300 at Ontario Power Generation's Darlington site, a four-unit Canadian program with a planning estimate around 21 billion Canadian dollars and a first unit targeting roughly 2030. One important caveat: much of what counts as "construction" so far is still early site works rather than full nuclear-island build, yet the direction is clear. After two decades of slide decks, the industry is finally pouring concrete. The right way to think about this is as an inflection from design-and-license to build-and-prove, where the open question shifts from "can it be approved?" to "can it be made affordable?"

Over the next two to five years, the energy system is under pressure from two directions at once. On the supply side, grids need firm, dispatchable, low-carbon generation to back up wind and solar when the weather doesn't cooperate. On the demand side, data centers and artificial-intelligence workloads are pushing electricity consumption upward, and nuclear regulators now cite that demand as an intensifying deployment driver. SMRs sit at the intersection, offering siting flexibility that reframes nuclear from a utility-scale public-works project into something closer to dedicated industrial infrastructure a single large customer might underwrite.

The deeper shift is in where competitive advantage sits. For conventional reactors, the moat was the ability to manage a massive one-off megaproject. For SMRs, the moat moves to the factory: whoever can manufacture standardized modules at volume, with a mature fuel supply chain behind them, wins. That relocates the bottleneck from construction sites to production lines and from engineering heroics to supply-chain discipline, and it makes serial order books, not single marquee plants, the metric that matters. Regulatory frameworks are being built in anticipation, including the United States' new technology-inclusive licensing rule known as 10 CFR Part 53, though the rule may still be in rulemaking and its real effect on conception-to-operation timelines is untested.

2. The Big Idea and How It Works

Think of the difference between a custom-built mansion and a fleet of identical prefab houses. A conventional gigawatt-scale reactor is the mansion: enormous, bespoke, built on site over many years, with every project relearning the last one's mistakes. An SMR aims to be the prefab house: major components fabricated in a controlled factory, trucked to site as modules, and bolted together. The old world treats each reactor as a singular engineering feat. The new world treats reactors as a manufactured product, where the value comes not from any one unit but from building the hundredth one faster and cheaper than the first. If it works, nuclear stops behaving like architecture and starts behaving like an assembly line.

Start with the defining feature. **Modularity.** Rather than constructing a reactor in place, manufacturers fabricate the major components as modules in a factory and ship them for on-site assembly, the same logic that lets factories outperform custom job sites on both speed and quality control. The benchmark to beat is the roughly eight-year average construction time for large reactors.

From a site operator's perspective, the flow is straightforward: order standardized modules, assemble them on a prepared site, then scale capacity by adding more identical units rather than designing a bigger one.

Two other concepts define the category. **Passive safety.** Many SMRs are engineered to shut down and cool themselves using natural forces, gravity, convection, and natural circulation, without external power or operator intervention. The reactor fails toward safe rather than relying on pumps and human action.

Fuel and reactor type. The landscape splits in two. Near-term designs (NuScale, the BWRX-300, Rolls-Royce's 470-megawatt unit) use familiar light-water cooling. Advanced designs use sodium, high-temperature gas, or molten salt. Several of these advanced designs, including the sodium-cooled Natrium, require HALEU, high-assay low-enriched uranium enriched above 5 percent but below 20 percent uranium-235, a fuel whose supply chain is not yet mature, making it a critical-path risk to roughly 2030 timelines.

3. Implications and Tensions

Products and companies. The economics flip from megaproject financing toward manufacturing scale. A smaller, factory-built unit is easier to site near a specific customer and easier to add incrementally, which advantages developers who can stand up production lines and threatens the assumption that nuclear must mean a single giant plant serving a regional grid. The plausible new customer is industrial: a data-center operator wanting firm, carbon-free power on a constrained timeline.

Work and workflows. The skills shift from on-site megaproject management toward factory production, module logistics, and standardized assembly. Quality control migrates from the field to the plant floor.

For something concrete and near-term: the BWRX-300 program at Darlington has firm milestones for its lead unit, targeting around 2030, and TerraPower's Natrium is advancing on its Wyoming site, the clearest current examples of Western projects crossing from concept toward construction, even as much of that early work is site preparation rather than nuclear-island build. Directionally, if serial production matures over the next several years, the medium-term picture is a repeatable product where each successive unit is built faster and cheaper than the last, turning order books into the key signal. But that outcome depends entirely on a manufacturing flywheel that no Western project has yet demonstrated.

Cost is the whole game. First-of-a-kind estimates run near $180 per megawatt-hour, while modeled mature ("Nth-of-a-kind") targets span roughly $52 to $100, with a U.S. Department of Energy analysis showing illustrative figures around $80 to $90. Every favorable number is a model output contingent on high-volume manufacturing that does not yet exist, and the OECD Nuclear Energy Agency frames cost competitiveness as a challenge to overcome, not an achieved result.

Speed versus proof. The two-to-five-year build timelines often cited are design targets, not demonstrated outcomes in Western contexts. No Western first-of-a-kind SMR has been completed, and early "construction" frequently means site prep rather than authorized nuclear-island work, so the schedule claim remains to be proven.

Hype versus deployment. More than 127 distinct SMR designs are tracked worldwide, yet only about seven were operational or under construction as of July 2025. The pipeline is enormous; the built fleet is tiny.

Promise versus permission. Passive safety raises the possibility of smaller emergency planning zones and siting closer to population, but that is regulator- and jurisdiction-dependent, a possibility rather than a settled property.

4. Hooks, Takeaways, and For the Nerds

A handful of lines that travel well in conversation:

- "The interesting bet with SMRs isn't the physics, it's whether you can build the hundredth one cheaply, like cars instead of cathedrals."
- "Russia and China already have these running. The West hasn't switched one on yet."
- "Every attractive cost number you'll see is a model output, not a real bill, contingent on factories that don't exist yet."
- "Early Western 'construction' mostly means site works. The first full nuclear builds will tell us if the two-to-five-year dream survives contact with reality."
- "There are over 120 designs on paper and roughly seven actually built or building. That gap is the whole story."

If you remember three things. SMRs reframe nuclear as a manufactured product, up to about 300 megawatts, factory-built in modules; the payoff is serial production, not any single plant; The decisive question is cost: first-unit estimates near $180/MWh must fall toward modeled $52–$100 targets through volume that hasn't been proven; Watch the lead Western builds (Natrium and the BWRX-300 at Darlington, both eyeing ~2030) and whether HALEU fuel supply arrives in time.

For the nerds. The light-water designs closest to construction represent evolutionary rather than revolutionary departures from existing reactor physics. NuScale's integral pressurized-water reactor (iPWR, where the steam generator and pressurizer sit inside the reactor pressure vessel) and GE Hitachi's BWRX-300 (a boiling-water design with a dramatically simplified containment) lean on decades of operational and supply-chain familiarity, with passive cooling that works through natural circulation rather than powered pumps.

The more ambitious designs introduce genuine novelty. Natrium's sodium-cooled fast reactor runs at near-atmospheric pressure and pairs with a molten-salt thermal-storage buffer, which lets electrical output vary independently of reactor power, a meaningful grid-services capability. The HTR-PM achieves passive safety through geometry itself: the fuel pebbles' low power density and high heat capacity mean the core physically cannot reach failure temperatures even in a complete loss-of-coolant scenario. But several of these advanced types need HALEU, and that supply chain remains immature enough to threaten 2030 schedules.

One genuinely open architectural question is waste. Smaller cores carry lower per-unit radioactive inventories, which sounds favorable, yet a distributed fleet of many small units could multiply decommissioning and waste-handling infrastructure across many sites. The source materials flag this as unresolved. The standing cautionary tale is NuScale's canceled UAMPS project, abandoned on ballooning cost estimates before construction, the clearest reminder that the economic flywheel is a hypothesis, not yet a result.

WHERE SOURCES DISAGREE

Will NOAK SMR costs actually fall to competitive levels?

One reading, drawn from DOE/INL techno-economic models and vendor targets, projects NOAK LCOEs as low as $52–$100/MWh once serial production matures. Another reading—from OECD-NEA and critical analyses—holds that SMRs forgo the economies of scale of large reactors and have not demonstrated that series-production learning will close the gap, with FOAK estimates near $180/MWh

Unresolved: every favorable figure is a model output contingent on high-volume manufacturing and supply chains that do not yet exist; the canceled NuScale UAMPS project (abandoned on ballooning cost estimates before construction) is the standing cautionary case

Is the "2–5 year" construction schedule achievable in Western regulatory environments?

The company and agency framing presents compressed factory-built schedules as a core benefit. The available evidence frames 2–5 years as an aspiration, with no completed Western first-of-a-kind build to validate it and a nuclear sector history of cost and schedule overruns

Unresolved: the first real tests (projects targeting ~2030) have not been completed

What is the actual nuclear construction status of the flagship Western projects?

One account presents TerraPower's Natrium and GE Hitachi's BWRX-300 as reactors "under construction." Another reading notes the sources distinguish non-nuclear site/early works from safety-related nuclear construction, and that firm dates and full-scope status are established mainly for lead units rather than entire multi-unit programs

Unresolved: the provided materials do not cleanly separate site-preparation works from authorized nuclear-island construction

Is the new U.S. licensing framework (10 CFR Part 53) finalized, and will it shorten real project timelines?

One reading treats Part 53 as a finalized, risk-informed, technology-inclusive rule streamlining advanced-reactor licensing. Another reading notes the rulemaking may still be in progress with

Unresolved: stakeholder comments, and that even if final, its practical effect on conception-to-operation timelines is untested. Unresolved because the framework is new and no developer has yet completed a project under it

Does a distributed fleet of small reactors increase or decrease total nuclear waste per kWh?

Some materials suggest smaller cores and lower per-unit inventories ease the waste picture; others note that many small, dispersed units may multiply decommissioning and waste-handling infrastructure

Unresolved: the source materials explicitly flag this as an open question rather than a settled empirical finding

Will HALEU fuel supply be ready in time for advanced non-LWR designs?

Advanced designs such as Natrium specifically require HALEU enriched between 5% and <20% U-235. The available evidence acknowledges this need but describes the supply chain as immature

Unresolved: the corpus flags the constraint without quantifying when supply will meet the ~2030 deployment timelines it threatens