Institute for Physical AI @ John Bailey Institute · The Charlot Lab & The Hiner Lab

ENERGY · NUCLEAR LOGISTICS

Technical Report TR–2026–39
Research / Position · Preprint v1
15 August 2026

ENERGY · NUCLEAR LOGISTICS

The Portable Core

Physical AI and the logistics of nuclear energy. The world is building the autonomy an unattended plant needs. It is not building the layer that lets the plant prove what it did.

The Charlot Lab & The Hiner Lab · Institute for Physical AI @ John Bailey Institute

Four-region literature sweep conducted 15 August 2026.

4 regions swept in native vocabulary | the premise was overturned | a null on our own seed terms | joint: two labs, one seam

Every embodied system ends in a joule supply chain, and portable nuclear power is the first energy technology that ships to its load rather than tethering the load to a grid. The dependency runs both ways: a sub-20 MWe plant cannot carry a 24/7 licensed crew, so the staffing model that works at 1 GW is the cost floor that kills 1 MW, and physical AI becomes not the customer of portable nuclear but its crew. This report opened from the assumption that the autonomy layer was unclaimed. **A four-region sweep in native vocabulary refutes that.** Korea operates a named Autonomous-Operation SMR Research Centre and targets standard design approval for i-SMR in 2028; China's nuclear safety regulator publishes reduced- and un-staffed intelligent operation as one of three named artificial-intelligence scenarios, including edge inference for autonomous response to reduce operator headcount; France frames NUWARD's passive safety as requiring neither electrical power nor human intervention. What none of them builds is a way for the plant to *demonstrate* what it did. Korea's digital twin monitors, diagnoses and predicts. EDF's simulates. Neither replays. China's artificial intelligence accelerates regulatory *review*; it does not make the artifact under review checkable. **That is the same structural gap the Institute's security series found for energy generally, power is telemetered everywhere and signed nowhere, and a mobile reactor crossing ports and jurisdictions is that problem at its hardest.** We state the thesis that the unattended plant is an evidence problem

> before it is an autonomy problem, identify the NUWARD Joint Early Review as the live multi-regulator venue where such an argument would land, and ship a runnable bench in which a regulator replays a filed shift against a deterministic twin, and discovers that the check is void when the plant authored its own log.

1. Why is the tether the problem?

Drones return to charge. Robots stop at the end of a cable. Forward bases burn trucked diesel at a burdened cost many times pump price once convoys and their protection are counted. Seabed sensors die when their batteries do. Edge datacentres queue for grid interconnects measured in years. Autonomy is discussed as a problem of perception and control; it is at least equally a problem of joules delivered at the point of action, and the point of action is rarely near a grid.

Portable nuclear inverts the relationship. Instead of moving the machine to the energy, the energy moves to the machine, and it does so as a logistics object, in denominations, on a ladder that now runs continuously from milliwatts to port scale.

2. What are the rungs of the ladder?

Rung	System	Scale and life	Status
mW–W	Betavoltaics	mW, decades	commercial, niche
W	Sr-90 seabed nodes	W-class, ~29 yr half-life	demonstration stage
10–100 kW	MARVEL	85 kWt / ~20 kWe	operation targeted late 2027
~1–5 MW	Kaleidos; eVinci	1.2 MWe/5 yr; 5 MWe/8 yr	entering test campaigns
10–20 MW	Installation class	≤20 MWe	nine sites named; demo targeted 2030
port scale	Barge-mounted SMR; **RITM-200S floating units**	175 MWt / 55 MWe per reactor, 7–10 yr between reloads	**four modernised floating power units under construction, two reactors each**

Table 1. One rung per class of machine, from a buried sensor to an automated port. The bottom row is the correction the sweep forced: a fleet of purpose-built floating units is a materially different fact from a single precedent vessel. VERIFIED, NATIVE-LANGUAGE SOURCES

3. What is the crew problem, and why was our first framing wrong?

Portable nuclear does not scale on human labour. A 1 MWe plant cannot carry a licensed operations crew, a security detachment and a maintenance staff around the clock. Every serious programme therefore assumes some combination of remote operation, autonomous control, robotic inspection and unattended security. That much this report asserted at the outset, and it holds.

What did not hold was the claim that follows from it. An earlier draft positioned autonomous microreactor operation as an open frontier that this Institute could enter. A sweep of four regions in their own technical vocabulary shows it is occupied.

Region	Autonomy posture	Twin
Korea	A named **Autonomous-Operation SMR Research Centre**; an eleven-organisation programme adding autonomous operation to SMR as source technology, with a systems-and-instrumentation team; operator minimisation explicit; i-SMR standard design approval targeted **2028**	synchronise real and virtual, then **monitor, diagnose, predict**
China	The nuclear safety regulator publishes **reduced- and un-staffed intelligent operation** as one of three named AI scenarios, edge inference for anomaly detection, fault prediction and **autonomous response** to cut operator headcount, alongside AI to **accelerate licensing review.** Inspection robots on fixed routes; autonomous aircraft with automatic battery exchange	—
France	Passive safety framed as requiring **neither electrical power nor human intervention**	**simulation** on a commercial modelling platform
Russia	not returned by the seed terms used	—

Table 2. The autonomy layer is claimed. VERIFIED

A methodological failure that nearly produced the opposite finding. The first two searches of this sweep used the terms an English-language reader reaches for (digital twin, unattended operation, autonomous operation) and returned hardware specifications only. That null would have *confirmed our own thesis*, which is precisely when it should be distrusted. Re-run with the industry's own vocabulary (reduced-staffing, intelligent operations and maintenance, instrumentation and control) the regulator's published position appeared immediately. **A null found on your own seed terms is a fact about your seeds.**

Correlation, stated next to the claim. Four regions are not four independent samples. Reactor engineering shares an international venue set and publishes substantially in English, so this is closer to one literature observed four times than to four draws. What genuinely differs between them is the *regulatory* posture, and that is where the rest of this report looks.

4. What survives?

Everyone is building the autonomy. This review did not locate work on making it checkable.

Korea's twin monitors and predicts (it is an operations tool, synchronised to the plant to anticipate failures. France's simulates) a design and modelling artifact. Neither is a bit-reproducible shadow that a third party could re-run against a claim the plant filed. China's use of AI in licensing accelerates the *review*; it does not change what is being reviewed into something a machine can check. And passive safety, which is real and important, removes the *need* for intervention without producing any *evidence* of what happened while nobody was there.

This is a shape the Institute has already priced. Its security series found that in general-purpose computing, power is telemetered ubiquitously and signed nowhere, and worked out what closing that costs: what a measurement plane can attest, whether a machine can demonstrate its own identity, and what a compliance gate does not constrain. An unattended reactor that changes ports, owners and jurisdictions is that same problem under the least forgiving conditions available.

Thesis: the unattended plant is an evidence problem before it is an autonomy problem.

5. What could a regulator check rather than trust?

Three artifacts follow, and each already exists in this Institute's work in a form that transfers.

Certificates that ride with the controller. Nuclear regulation runs on demonstrable safety cases. A controller carrying a stability proof checked on the device offers a regulator something categorically different from an empirical performance claim: an artifact that can be evaluated rather than believed. This is a licensing posture, not a benchmark.

A deterministic, replayable twin. Distinct from every twin the sweep located. The property that matters is not fidelity but *reproducibility*: given the same commanded sequence, the twin produces the same output, so a filed claim can be re-run rather than accepted.

Generation receipts, with the constraint stated in the spec, not discovered later. Extending verifiable energy receipts from computation to production is the obvious move for an asset that changes jurisdictions. It is bounded by a result this Institute has already published: **a receipt computed by the host is the host's account of itself**. Signing it changes its provenance, not its truth. Any generation-receipt specification must therefore name who measures, and what makes that measurement unforgeable by the plant's own software, before it is a specification at all.

Figure 1, the replay bench, runnable. A plant files a claim about a shift. The reader is a regulator who cannot visit the site and cannot trust its software, holding a deterministic twin.

Log source	Shift	Hours flagged	Discrepancy	Verdict
independent meter	nothing hidden	0	0.00 MWh	passes
independent meter	undeclared derate	6	1.30 MWh	discrepancy found

the plant's own software	undeclared derate	0	0.00 MWh	VOID

The third row is the report's point. The plant files exactly what the twin expects, the replay agrees perfectly, and the derate is still there. Widening the twin's fidelity band past the depth of the shortfall reproduces the same silence from the other direction. Run it. MODELLED; NO FIGURE HERE IS A MEASUREMENT OF A REAL PLANT

6. Why is this a two-lab report?

The problem divides along a real seam. **The plant boundary** (matching a flat core to a bursty embodied load through storage, compute as dispatchable load, and thermal offtake as a second port) is impedance matching at site scale, and it is where the economics are decided. **The evidence layer** (certificates, replayable twins, receipts rooted outside the host) is what makes the boundary's instrumentation mean anything to a third party.

Either half alone is what the sweep already found being built somewhere. The pairing is the contribution.

7. What would sharpen this?

Transportable-reactor licensing is less mature than stationary microreactor licensing, and a mobile core carries security and safeguards questions that a fixed one does not. Fuel supply is constrained. Public acceptance near ports is untested at scale. Most published microreactor economics are models rather than measurements and are marked as such throughout.

And the field has a graveyard. The 1960s portable military reactor programmes died on economics and maintenance, which is exactly the failure mode autonomy must be *shown* to fix rather than assumed to. This report takes the position that a demonstration of checkability is a stronger response to that history than a demonstration of capability.

The sweep covered four regions on stated seed terms, one of which returned nothing on the autonomy question. That is recorded as not returned by those seeds, which is not the same as not existing.

8. Conclusions

Portable nuclear turns energy into a logistics object, and physical AI turns the plant into an unattended machine. The second half is being built in at least three countries, with a named research centre, a regulator's published position, and a design-approval target inside three years. The half that is not being built is the one that lets an unattended plant prove what it did to someone who was not there and has no reason to trust it.

That is the opening, it is narrow, and it is the one this Institute is equipped for. The report's own opening assumption, that autonomy was the unclaimed ground, did not survive its own literature sweep, and is corrected here at the size it was claimed.

9. The forcing function

This report finds a layer unbuilt. That is a schedule, not a verdict, and it is worth writing out as one.

What is missing	Why it is missing	What builds it	What becomes possible
Evidence from an unattended plant	Nobody was present, and the plant's own software is the only witness	A measurement rooted where the plant cannot author it, plus a **replayable** twin rather than a monitoring one	A plant that can be licensed on what a regulator checks rather than on who vouches for it
Generation receipts that survive a change of jurisdiction	Receipts today are accounting artifacts computed by the asset	Naming the measurement root in the specification's first line	Energy as a genuinely portable asset, whose history travels with it
A boundary that absorbs a flat core	Curtailment and unmet demand are different failures and are usually conflated	Sizing storage against timing and sinks against capacity, with dispatch order specified	The small rungs of the ladder becoming economic, which is what puts power where machines actually work

Table 3. Three engineering tasks, no new physics among them.

The 1960s portable reactors were **impossible** to run economically because they demanded continuous skilled attention. Three countries have since made unattended operation *probable*, a research centre, a regulator's published position, a design-approval target inside three years. What makes it **possible** is the half this review did not locate work on: the ability of a machine standing alone in the world to prove what it did. That is an engineering-physics problem about where a measurement is rooted, and it is the one this Institute is equipped to take.

References

1. Four-region literature sweep, 15 August 2026: Korean, Chinese, French and Russian sources in native vocabulary. Per-region findings and seed terms in the working spine. VERIFIED; CORRELATED SAMPLE, SEE §3
2. Institute for Physical AI @ JBI, TR-2026-36, *Energy Observability in Embodied Systems*; TR-2026-37, *Can a Machine Prove It Is Itself?*; TR-2026-38, *A Certificate That Bounds Danger Does Not Bound Waste*. INSTITUTE PUBLICATIONS
3. Institute for Physical AI @ JBI, TR-2026-07 (provable by construction), TR-2026-20 (the matching principle), TR-2026-04 (microfluidic co-optimisation), TR-2026-05 (energy-native compute), TR-2026-17 (managed low-altitude corridors). INSTITUTE PUBLICATIONS
4. Multi-regulator joint early review of a European SMR design: pilot phase closure 2023, phase 2 summary December 2025, phase 3 opened 21 January 2026 with eight participating regulators. VERIFIED; PRIMARY DOCUMENTS PUBLIC

▶ Run the bench

Institute for Physical AI @ John Bailey Institute · The Charlot Lab & The Hiner Lab

TR-2026-39 · Preprint v1 · 15 August 2026