Permanently Shadowed Regions as a Shared Compute / ISRU Hub

An architecture position paper. Separate from, and more speculative than, the low-grade waste-heat-offset paper. Its claims stand on their own and should be evaluated independently.

Version 0.2 (2026-07-01). This paper argues a *siting* thesis; the quantitative energy-offset kernel it leans on is the companion `WASTE-HEAT-OFFSET.md`. Where that paper makes a narrow, Second-Law-safe claim, this one makes a broader systems argument that carries more assumptions, called out explicitly.

Thesis

A permanently shadowed region (PSR) is, at the same time, **(1) the best lunar location to reject compute waste heat** and **(2) where the water resource is**. Those two facts are usually discussed in separate literatures. Put together, they suggest co-locating a surface compute facility with a water-mining ISRU plant inside or adjacent to a PSR, and cascading the compute's rejected heat into ice sublimation before dumping the remainder to the cryogenic PSR sky. The PSR turns a data center's largest liability (heat rejection) into a shared asset, and amortizes the one hard thing both systems need: power delivered into permanent shadow.

Claim 1: A PSR is a superior compute heat-rejection environment

Radiator performance is set by what the radiator *sees*. On the sunlit equatorial surface a radiator absorbs direct and reflected sunlight and exchanges with ~250-400 K terrain, so its effective sink is warm and its area is large. Inside a PSR there is **no direct solar load** and the surrounding terrain sits at ~40-110 K, so a radiator can run cold and small, and a given compute load rejects its heat with far less radiator mass. Spacecraft thermal practice already treats a cold, sunless view as the ideal rejection condition (see the spacecraft-thermal-control literature); a PSR is the most accessible place in cislunar space that offers it on the ground. This is a genuine, under-exploited siting advantage for compute, independent of ISRU.

Claim 2: The waste heat then lands exactly where ISRU needs low-grade heat

Per the companion paper, the only ISRU heat demand a ~315-350 K waste stream can serve is the PSR water route's sublimation chain (~273 K target), and it serves all of it (~1.8 kWh/kg O2 after the exchanger pinch and effectiveness, ~11% of the route). Because the resource is *in the PSR*, co-location removes the transport problem that would otherwise kill the idea: the heat source, the low-grade sink (icy regolith), and the final cold sink (PSR sky) are all in the same place. A ~12 kW compute load covers a 50 t/yr water plant; on energy grounds a 100 kW facility could serve ~500 t/yr, though delivering heat into a granular icy bed is conduction-limited, so that is an upper bound, not a design. The cascade is: compute silicon → coolant loop → ice sublimation → residual to radiators.

Claim 3: Co-location amortizes the power-into-shadow problem

The decisive unsolved constraint for *both* systems is the same: a PSR receives no sunlight, so power must come from fission (a reactor sited on a sunlit rim or in the PSR) or be beamed in (demonstrated electrical-to-electrical efficiency only ~11.5% at 10 m). Whichever is chosen, a co-located compute + ISRU hub pays for that infrastructure once and shares it, rather than two separate programs each solving permanent-shadow power independently. The compute load and the ISRU load are also complementary in time: compute is a steady baseload, ISRU can be throttled, which smooths the demand on a reactor.

Quantified benefit, and the probability it is realized

The two benefits are distinct and should not be summed; conflating them oversells the cascade (`python -m lpem --benefit`, reference: a 50 t/yr PSR water plant + co-located compute):

- **Cascade benefit** (reuse compute heat in ISRU): saves ~**2.3 t** of landed reactor mass (the low-grade heat offset, avoided fission power), against ~**1 t** of heat-integration hardware (exchanger, transport loop, dust mitigation).
- **Siting benefit** (put the compute in the PSR for its cold sink): saves radiator mass, scale-dependent. Under an explicit radiator energy balance (IR emission minus absorbed solar minus environmental IR, with a sky view factor for a competently oriented vertical two-sided panel, F__sky nominal 0.5), a PSR saves, over the feasible sampled cases, a median ~**10 t of radiator per MW of compute** (IQR ~7-16 t/MW). An earlier version of this paper reported ~29 t/MW and claimed ~30% of sunlit designs could not reject at all; both figures were artifacts of implicitly modeling a horizontal panel staring at warm terrain with zero sky view. With realistic vertical-panel geometry, only ~0.1% of sampled sunlit designs truly cannot reject at 330 K (net flux non-positive at any area) and a further ~0.4% can reject only with a prohibitive area (more than 10x the PSR panel); the corrected advantage is the ~10 t/MW mass saving, not widespread sunlit infeasibility. At data-center scale this still exceeds the cascade and remains the actual driver of co-location. It is a *siting* benefit, not a cascade benefit: a PSR already offers a cheap radiative sink, so the cascade itself does not save radiator mass. (Model: `src/lpem/benefit.py`, `net_rejection_wm2`; parameters, emissivity, radiator temperature, absorbed solar, environmental IR, sky view factor, are explicit and tunable.)

Estimating the probability of the benefit. Rather than assert subjective probabilities (P that lunar-surface compute exists, etc.), we compute the **break-even joint probability** the enabling chain must clear for the cascade hardware to pay for itself: P* = cost / benefit, ~**44%** at nominal and ~**50%** once propagated (the ratio is right-skewed, so the nominal under-reports it). Two readings follow:

- *Conditional on co-location already happening* (compute and a water plant both sited at the PSR for their own reasons), the cascade is worthwhile in ~**78%** of trials (the integration probability clears the propagated break-even), so it is a favorable, positive-expected-value add-on, though a closer call than the earlier ~85% estimate (the cascade prize shrank once the offset model gained a pinch and effectiveness). **Design it in.**
- *As a standalone speculative bet*, the full enabling chain (surface compute exists, co-located, water route pursued, integration works) under wide illustrative priors is only ~9%, well below the break-even, so its expected value is negative (~-0.8 t). **Do not justify co-location by the cascade alone.**

The decision structure is therefore clear: **co-location is justified (if at all) by the compute siting economics (roughly ten tonnes per MW of radiator mass and shared power-into-shadow infrastructure) and the ISRU heat cascade is a cheap, sensible bonus to capture once you are already there, not a reason to go.** The break-even model (`src/lpem/benefit.py`) makes every assumption explicit and tunable.

What this is not, and what would sink it

This is a siting argument, not an engineering design, and it rests on assumptions a skeptic should press:

- **Why surface compute at all?** The thesis is empty unless there is an independent reason to run compute on the lunar surface (autonomy and real-time ISRU/robotics control, proximity to a future lunar data economy, or latency-tolerant batch compute that wants a free cold sink). If all useful compute lives on Earth or in orbit, there is no waste heat at the PSR to cascade. This is the load-bearing assumption.
- **Power into shadow is genuinely unsolved.** Co-location amortizes it but does not solve it; if neither in-PSR fission siting nor efficient beaming matures, the whole hub is blocked, and so is standalone PSR water mining.
- **Heat transport over the last tens of meters** (racks to the working face) still costs mass and pumping power, and the regolith-side heat exchanger operates in abrasive dust at cryogenic temperatures with no flight heritage.
- **The energy prize is modest** (~11% of one route); the real argument is shared infrastructure and a free, ideally-located heat sink, not a large kWh saving.

- **Dust on radiators** in a worked mining environment degrades the very rejection advantage Claim 1 depends on.

What would make this a real study

A defensible follow-on needs: a concrete surface-compute demand case (kW and why it is on the Moon); a co-located mass model (reactor + radiators + compute + mining, shared vs separate) showing the amortization quantitatively; a heat-transport design from racks to the sublimation face; and a comparison against the obvious alternative, solar-thermal process heat from a sunlit rim with uphill resource haul. The `lpem` energy and power/landed-mass models are the natural backbone for the mass-amortization piece.

Relationship to the other papers

- `WASTE-HEAT-OFFSET.md`, the narrow, quantified, Second-Law-safe energy claim this paper builds on. That paper stands without this one.
- `MANUSCRIPT.md`, the underlying common-basis energy model and the route rankings.