

The Cloud-Grid Settlement: Provenance and Horizon

How this was drawn up, and what an achievement of this kind would make possible

*A companion to the term sheet, for circulation alongside it. Receipts discipline applies throughout: load-bearing claims are tagged **[provable]** (named, dated source), **[derived]** (computed from stated inputs), or **[inference]** (a structural judgment, stated as such). Provable carries the argument; nothing speculative is allowed to read as settled.*

Pairs with cloud-grid-termsheet_v2026-06-22c, which reached its current form across three rounds of adversarial review: the instrument was re-based from royalty to systemic-risk fee; firm status settled at scoped co-governance with a recusal bar (Option C) after voice-without-vote was found to leave a fiscal hold-up; and the continuity guarantee was redesigned from pre-clearance of a sovereign act (unsignable) to continuity-of-capability (escrow, portability, stand-by capacity) plus automatic pre-costed reversal. The review record is part of the provenance: the instrument is sound, the early contradictions are repaired rather than relabeled, and the surviving hard question is feasibility — whether a single calibration satisfies all parties — which the term sheet now begins to test with indicative figures rather than defer.

This note exists because a serious proposal should be able to account for itself. A reader handed a term sheet is entitled to two things the term sheet itself does not provide: where it came from, and where it leads. The first is a question of method — whether the settlement was reasoned into existence from evidence or asserted from ambition. The second is a question of stakes — whether an achievement of this kind is a narrow fix or the opening of a larger possibility. What follows answers both, in that order, and holds itself to the same standard of proof the term sheet does.

How this was made

A word on method, because it is part of what is being offered rather than something to apologize for or conceal. This is open, labor-grounded research of a kind that barely exists yet — an early attempt at it — produced by one person working from inside the field rather than studying it from outside. It was built in a compressed window using AI research tools. That fact is not incidental and is not hidden: the scope here — a systemic-risk instrument, multi-sovereign governance precedent, a labor and safety-net analysis, and three rounds of adversarial review — is ordinarily the work of a team over weeks, and the tooling is what let one person cover that ground in the time available. The usage is a function of time and scale, nothing more.

What keeps that honest rather than hollow is the discipline wrapped around it. Every load-bearing claim is labeled for what it is — provable, derived, or inference. The weak points are named in writing rather than smoothed over. The structure was put through repeated adversarial review specifically to find what would break, and what broke was cut and recorded. The method, in other words, was designed so that the work can be *checked* rather than *trusted* — which is the right posture for any document built with these tools and asking to be taken seriously, since such a document sits inside exactly the suspicion it has to answer. The response was to over-document, not to ask for the benefit of the doubt. The receipts discipline is the point, and it travels with the work.

The starting line was a single question — what happens, and who bears it, when the infrastructure the world now runs on fails all at once — and everything below is the trail from that question to a settlement, with the labor and safety-net grounding that gives it a public-interest core rather than a purely technical one. The deeper soil it grew from (long-running work on the trades, on community health workers, and on the receipts method itself) is real but kept offstage here; this account starts where the workstream started.

Part I — Provenance: how and why this was drawn up

The settlement was not designed top-down from a desired conclusion. It was assembled from a single empirical finding that survived deliberate attempts to destroy it, and then built outward only as far as the evidence would carry. The sequence matters, because it is the sequence that makes the result defensible rather than aspirational.

The finding that started it. The work began with a structural claim about the hyperscale cloud: that its correlated catastrophic failure has no risk-transfer mechanism — no private insurance market that can price it, no statutory backstop that absorbs it, and, on the one real precedent, no public catch either. That claim was put through four successive adversarial reviews specifically to find what was overstated, and most of the original thesis did not survive. **[provable]** What was killed is part of the receipt: the claim that “the public is the de facto insurer,” any headline catastrophe dollar figure, a dependency multiplier, “exceeds all reinsurance,” and “unallocated” as a dramatic frame were all struck as unprovable or imprecise. What survived is narrow and solid: concentration of roughly two-thirds of global cloud in three US firms; the structural uninsurability of the correlated tail (correlation defeats pooling, cyber exclusions, the terrorism-certification gate in existing federal backstop law, open GAO recommendations); and *loss-lies-where-it-falls*, demonstrated by the July 2024 CrowdStrike event, in which roughly 8.5 million machines failed, losses ran to the billions, only a fraction was insurable, vendor liability was capped at “fees paid,” and government — among the victims — socialized nothing. **[provable]** The discipline of that strip-down is the foundation: the settlement rests only on what the reviews could not take away.

The exposure inversion that gave it strategic shape. A parallel analysis mapped how the dependence is distributed, on two distinct axes kept rigorously separate — *structural* exposure (the architecture of dependence, documentable) versus *systemic* exposure (whether a shock cascades, a directional judgment only). **[provable for the structure; inference for the cascade]** The result was an inversion: US adversaries are largely insulated (China substantially domestic by design; Russia and Iran cut off by force), while US allies are maximally wired in (Europe, Israel, the Gulf, India). The grid is a chokepoint held over friends, not adversaries. That inversion is what turns a domestic risk question into a settlement question — it identifies who the counterparties are and why they are exposed.

The grounding that made it more than a finance proposal. Separate research lines established the domestic substance the settlement would have to serve. Investigation into data-center construction under organized-labor agreements found, among other things, that no dedicated study of the question even exists — an open finding in itself — and documented the building-trades posture toward the buildout, including the NABTU–OpenAI memorandum and the internal labor split between trades that see a generational opportunity and a contingent calling for moratoria.

[provable] Verification of public-opinion figures on the buildout corrected a widely-circulated misattribution and established the real numbers (roughly 71% opposition to a local data center in one major poll; 57% in another), separating the provable from the inflated. **[provable]** And a methodological line on household-welfare measurement established how a defensible composite index is built to survive hostile review — Alkire-Foster aggregation with explicit weights, a published dashboard, formal sensitivity analysis. **[provable]** These are not decorations on a financial instrument; they are the reason the instrument has a public-interest core rather than a purely technical one.

The correction that fixed the instrument. An early version of the levy was anchored to resource royalties — oil-and-gas benchmarks. That was a category error, and naming it is part of the receipt:

resource royalties rest on public ownership of an extracted resource, and there is no public ownership of privately built compute. **[provable]** The instrument was re-based onto the family of mechanisms whose animating logic actually fits the asset — financial-stability surcharges (deposit-insurance assessments, a resolution fund, a systemic-importance surcharge), which exist precisely for a correlated tail backed by a de facto public guarantee with no private market. **[provable]** That re-basing is the single decision that lets the number survive a hostile economist, and it is why the settlement is a *systemic-risk fee*, not a tax or a royalty.

The synthesis. Only after those pieces were each independently grounded were they assembled into the settlement: an outward term sheet proposing, to allied states and to the firms, a priced and governed and resilient grid — and an inward grounding (local rent capture, the workforce baselines, a community health corps) that builds the domestic durability the outward terms require. The connection between them is the load-bearing architecture: a counterparty signs only if it believes the United States can sustain its side across election cycles, and a durable domestic majority behind the deal is what makes that believable. The grounding underwrites the term sheet the way a balance sheet underwrites a guarantee. **[inference]**

That is the whole provenance: a finding hardened by adversarial review, a strategic shape from the exposure inversion, a public-interest core from the labor and welfare work, an instrument corrected onto the right precedent, and only then a settlement. The why is simply stated — the finding identified a real, unowned risk and a real, divided majority, and no existing instrument addressed either. The settlement is the constructed answer, built no larger than the evidence permits.

Part II — Horizon: what lies downstream of an achievement of this kind

The term sheet is itself a horizon — a destination, not a near-term task, gated on a domestic political precondition stated frankly elsewhere. This section looks past it: if the settlement were achieved, what would it open? The honest answer requires the same discipline as the rest. Each possibility below is tagged for how far the evidence reaches, and the further downstream it sits, the more speculative it is allowed to be only if it says so.

The risk-transfer market it was built to diagnose would finally form. The inquiry began from the absence of any mechanism to price the correlated tail. A standing, pre-funded backstop priced like deposit insurance is exactly what lets a private residual market form on top of it: once the catastrophic layer is absorbed and priced publicly, the remainder becomes underwritable. This is how terrorism reinsurance was made writable after a public backstop was created — the public layer de-risks the tail enough that private capital returns to the rest. **[inference, precedent-grounded]** The settlement would, in other words, resolve the very finding that started it.

A universal-access floor would become fundable rather than aspirational. A standing fund with a contribution mechanism is the missing vehicle beneath every existing universal-access commitment — the connectivity targets, the affordability thresholds, the digital-public-infrastructure templates already proven at national scale. The telecom universal-service model shows both the mechanism and its failure mode (a shrinking contribution base inflates the rate until it collapses); a fund assessed on a broad and growing hyperscaler base is structurally better positioned than its precedents. **[derived, precedent-grounded]** The floor stops being a menu of standards and becomes a line item.

The institutional form would transfer to other concentrated chokepoints. A systemic-risk fee paired with capped, weighted-but-protected governance is not specific to cloud. It is a reusable form for any privately held, systemically important infrastructure where concentration has outrun

governance — undersea cables, payment rails, satellite constellations, and eventually the foundation models themselves. The bank-surcharge framework it borrows from spread across the entire global banking system once it existed in one place. **[inference]** The settlement would be a template, not a one-off.

The bloc inversion would be repaired rather than merely named. The exposure analysis identified the strategic vulnerability — a chokepoint held over allies — as the thing most worth fixing. A settlement that binds those allies into shared governance and guaranteed continuity converts chokepoint-over-friends into governed shared infrastructure; and because accession is gated by acceptance of binding terms rather than by alignment, the destination is not an allied bloc but a commons open in principle to any state, adversaries included, that accepts the same discipline — a rival reducing its own exposure by joining rather than by building a competing stack. The precedent for the founding move is the satellite consortium that bound allied states into a US-dominant-but-capped structure they accepted because the alternative was more expensive than buying in; the precedent for the open door is the way the major arms-control and monetary institutions admitted rivals on condition of accepting their constraints. **[inference, precedent-grounded]**

The domestic grounding would outlive the deal as permanent institutions. The workforce baselines and the community health corps that underwrite the settlement's credibility are not consumed by it. A standing skilled workforce built to a diversity-and-parity standard, and a community health corps funded in part from the dividend, would persist as institutions after the settlement that called them into being. **[inference]** The grounding is the part that compounds locally regardless of the horizon's fate.

A working vehicle for AI-rent capture generally would exist. Proposals to capture AI rents into a sovereign fund currently lack an instrument — they name a destination without a rail to reach it. A functioning dividend mechanism with a defined assessment base, a governed fund, and a distribution rule is that rail, available for any future rent-capture commitment to run on. **[inference]**

Democratic legitimacy for AI governance would have a standing home. The most powerful infrastructure of the period is currently governed ad hoc, by the firms and a handful of states. A standing multilateral body with a basic-vote floor that preserves the smallest parties' weight is the institutional answer to “who governs this, and by what right.” **[inference]** This is the furthest-downstream and most speculative item, and it is named as such.

A closing note on method

The two parts of this note are deliberately asymmetric in their burden of proof. The provenance is mostly **[provable]** — it is a record of what was established and what was discarded, and it is meant to survive verification. The horizon is mostly **[inference]** — it is a reasoned account of what an achievement would open, grounded in real precedents but speculative by nature, and it says so at every step. Holding those two registers apart is the whole method. A proposal that lets its hopes for the horizon leak backward into its claims about the finding is how ambition corrupts evidence; keeping the provenance hard and the horizon honestly soft is what allows the settlement to aspire to stature without overreaching its receipts.

The precondition that gates all of it is stated plainly in the settlement itself and is not softened here: the horizon is not executable under present conditions, and the grounding is the near-term work that builds the capacity to reach it. This note does not pretend otherwise. Its purpose is narrower and, within its limits, complete — to show that the settlement was drawn up from evidence rather than appetite, and that what it points toward is worth the building.