

A Cloud-Grid Settlement

One worked proposal for governing concentrated cloud and AI infrastructure

Public summary · Version 2026-09-19a · Instrument dated June 22–24, 2026

What this is

The world now runs on a handful of cloud providers concentrated in one country. Their simultaneous failure is something no insurance market can price and no government has agreed to absorb, and the dependence falls hardest on that country's closest allies. This is one worked answer to the question that raises.

It is a settlement: a treaty-style instrument that prices the systemic risk through a fee, governs the infrastructure jointly rather than privately, guarantees that a dependent country keeps operating even if its service is cut, and dedicates part of the proceeds to a floor of universal access. It is written as indicative heads of terms — what each party would agree to, and what each would get — pitched at the level of architecture, with the detailed calibration named and deferred rather than faked.

It is non-binding and pre-decisional. It is offered to be tested and argued with, and if someone can do better, the sooner the better.

Why it exists

A great many serious people now argue that concentrated cloud and AI capacity should be treated as a commons, a public utility, or critical infrastructure. Far fewer have said how such a thing would be paid for, or who would govern it, or what a state that joined would actually receive in return. This instrument is an attempt at that missing part: the machinery rather than the frame.

The method was to work several steps ahead of current events, including the severe cases, to see what would eventually have to be addressed, and then to work back to terms the parties could negotiate from where they stand now. The result is meant to look like the end product of a difficult negotiation that is achievable at scale, rather than a wish list.

The four moving parts

A price on the risk. A fee on providers above a size threshold, built the way large banks are made to pay into a fund against their own failure. Concentration creates a risk the public currently absorbs for nothing. The fee prices it, and the fund it builds is what pays for the guarantees below.

Shared governance. The rules of service written jointly by the states and the providers in the settlement, rather than by a few firms alone, with the controlling country's weight capped and declining and a floor of voting weight for the smallest members.

A continuity guarantee. A country that depends on this infrastructure keeps running even if its service is cut — delivered by architecture (escrow, portability, stand-by capacity) rather than by anyone's permission.

A floor of universal access. Part of the fund pays for access for countries that would otherwise be priced out entirely, ring-fenced as a first call so that a crisis elsewhere cannot defund it.

The door

Any state can join, rivals included, on the same binding terms. Accession is gated by acceptance of the constraints, not by alignment. This is the difference between a defensive bloc and a settlement: a rival reduces its own exposure, and the system's, by accepting the same terms as everyone else rather than by building a competing stack.

The documents

Five documents are published together, in reading order:

- This summary — the whole instrument in brief.
- The Term Sheet (v2026-06-22c) — the instrument itself: the fee, the governance, the continuity guarantee, the access floor, and what each party gives and receives.
- Provenance and Horizon (v2026-06-24a) — how the instrument was reasoned from evidence rather than asserted from ambition, including a record of what earlier review killed, and what an achievement of this kind would open downstream. Published as a public edition, with an internal appendix withheld.
- A Prospective Settlement for the AI Grid (v2026-06-24e) — a plain-language introduction through a single case, for a first-time reader.
- Sourcing and Receipts Index (public edition, v2026-09-19a) — every load-bearing outward claim tagged provable, derived, or inference, with the weak points named and three corrections flagged.

Status and method

Everything here is a draft and marked as such. The instrument was written June 22–24, 2026, circulated privately to a short list on August 6, 2026, and published and deposited in September 2026. The June dates are the record, and nothing in the original documents has been revised toward later events. Three load-bearing figures were corrected before the documents were first circulated, and those corrections are recorded openly in the Sourcing Index rather than made silently in the text. This summary is the one new document in the set, written in September 2026 for readers meeting the work for the first time.

The work was produced by one person in a compressed window, using AI research tools (Anthropic's Claude) for source gathering, drafting, and internal consistency checks, under a discipline designed so that the result can be checked rather than trusted: every load-bearing claim labeled for what it is, the

weak points named in writing, and the open problems listed rather than smoothed over. The author reviewed the material and stands behind it.

No outside expert has reviewed this work. Earlier rounds described as adversarial review were internal consistency checks run through the same system that produced the drafts, which is a useful discipline but is not independent review.

DRAFT