

CLOUD-GRID SETTLEMENT — INDICATIVE HEADS OF TERMS

The world now runs on a handful of cloud providers concentrated in one country, whose simultaneous failure no market can insure and no government has agreed to absorb — and the dependence falls hardest on that country's closest allies. What follows is one worked answer to the question that raises: a settlement that prices the systemic risk, governs the asset jointly, guarantees that a dependent state keeps operating even if its service is cut, and returns part of the proceeds as a universal-access floor. It is written as a set of indicative 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, offered to be tested and argued with. A full note on status, method, and version sits at the end, where it belongs; the terms come first.

Recitals

A. A majority of the world's cloud infrastructure runs on a small number of providers domiciled in one jurisdiction. The concentration is a matter of public record.

B. The correlated catastrophic failure of that infrastructure has no private insurance market able to price it and no statutory mechanism to absorb it. On the one large-scale precedent to date, the loss was borne where it fell, including by public bodies, and was not transferred or socialized by design.

C. The dependence is distributed unevenly: it falls most heavily on the states most closely aligned with the provider jurisdiction. Shared infrastructure of this importance, left ungoverned, is a shared exposure.

D. This document sets out indicative terms under which the provider jurisdiction, allied states, and the providers themselves would place that infrastructure under a priced, governed, and resilient settlement — pricing the systemic risk, co-governing the resilience of the asset, guaranteeing continuity of capability, and dedicating part of the proceeds to a universal-access floor. ***The founding cohort is the provider jurisdiction and its allies, where the dependence is deepest and the trust to begin already exists; but the settlement is open in principle to any state that accepts its terms*** (the Assessment, the Continuity-of-Capability Guarantee, the non-weaponization rules, and the governance discipline). Accession is gated by acceptance of the constraints, not by alignment — the horizon is a governed commons, not a bloc, and a non-allied state may join only by accepting the same binding terms as everyone else, on the model of how the major arms-control and monetary institutions were built.

E. The terms are pitched at the level of architecture. Detailed calibration — provider-by-provider rates, jurisdiction-by-jurisdiction schedules, stress-tested fund sizing, the legal form of cross-border liability channeling, and the engineering of the continuity circuit-breaker — is deliberately left to the work that would follow accession (Part K). Where a number is needed to make a mechanism testable, it appears below marked **illustrative** and carries no commitment.

F. **On the controlling jurisdiction's reason to enter.** A settlement of this kind is not what a dominant holder of the asset does at the height of unchallenged control; it is what a holder does when it judges that undivided control is eroding — allies building sovereign alternatives, the political cost of the chokepoint rising — and that durable cohesion is worth more than formal sole control. These terms are addressed to a future government of the provider jurisdiction acting on that

judgment; they do not assume the incumbent's incentives, and they do not attempt to supply the political precondition. The fuller account of who signs, and why, is held in the companion notes and is deliberately off the face of this document. What follows assumes a willing controlling jurisdiction and specifies what it would agree to — and, in Part D, has been redesigned specifically so that what it is asked to agree to is something a sovereign realistically can.

Part A — Defined Terms

Settlement — the arrangement constituted by these terms once agreed.

Council — the governing body: acceding States as voting Parties; Covered Providers as scoped co-governors of the resilience layer (Parts C, E).

Party — an acceding State holding a seat and full vote.

Covered Provider — a designated systemically important cloud provider: an assessed entity co-governing the resilience layer only, under a binding Recusal Bar (Part C).

Assessment — the systemic-risk fee paid by each Covered Provider (Part B).

Assessment Base — the revenue measure on which the Assessment is levied (clause 9).

Fund — the ring-fenced fund into which Assessments are paid and from which the Access Floor (first call) and Coverage are paid (clause 11).

Resilience Standards — the technical continuity, redundancy, and portability requirements; the **resilience layer** is the set of decisions concerning these Standards.

Continuity-of-Capability Guarantee — the structural assurance that a Party retains operational capability notwithstanding a service action (Part D).

Schedule Y — the enumerated, exhaustive list of lawful grounds on which service may be interrupted.

Access Floor — the universal-access entitlement funded from the Fund as a first call (Part F).

Reserved Matters — decisions requiring a supermajority and protected from simple-majority change (clause 19).

Recusal Bar — the matters on which a Covered Provider may not vote, deliberate, or sit (clause 13).

Adjudicator — the neutral body resolving disputes, post-action review, and Assessment appeals (clause 20; identity an Open Item).

Part B — The Instrument

6. The Assessment. Each Covered Provider pays an annual Assessment into the Fund. It is a systemic-risk fee — not a tax on profit, not a royalty on output. Each component is drawn from the mechanism that actually calibrates it, and no further: (a) **Form** — basis points on the Assessment Base, paid into a pool, on the model of deposit-insurance assessments (a fee into a pool, not capital held against a firm's own losses). (b) **Absolute rate** — set by the Fund target: the basis points are whatever is required to build the Fund to its target over the glidepath, on the model of a pre-funded resolution fund. The rate is a function of the target, never of any firm's profit. (c) **Relative share** — how the aggregate is *apportioned among providers* is set by a five-indicator systemic-importance score (size, interconnectedness, substitutability, complexity, cross-jurisdictional footprint), adapted from the methodology used to rank systemically important banks. The score allocates shares; it does not size the fee. (The bank surcharge is a capital requirement, not a fee; only its ranking logic transfers, and only to apportion.)

7. Entry rate. The entry band is **1 to 5 basis points of the Assessment Base** — on the present aggregate base, on the order of **\$25–130 million per year**. This is a starting rate for the accession period, not a ceiling and not the steady state. It rises to whatever the Fund target requires (clause 8). The low entry is a function of phasing, not of the Fund's eventual size.

8. Escalation, and what the Fund is sized to — with illustrative figures. The rate rises from entry to steady state along one governed path: the Fund target. The Fund is **not** sized to absorb a full correlated-catastrophe loss outright — no pre-funded vehicle of plausible size could, and these terms do not pretend otherwise. It is sized to a defined **first-loss layer**: the initial, mutualizable band of a correlated event, which is the band that seeds pricing and an attachment point and so makes the residual privately insurable (clause 14, Part J). *To make the mechanism testable rather than abstract, and marked illustrative* — carrying no commitment and subject entirely to the calibration in Part K: a first-loss layer on the order of **\$15–40 billion**, reached over a **10–15 year** glidepath, would imply a steady-state Assessment on the order of **30–80 basis points** of the present Base (roughly **\$0.8–2 billion per year** before growth in the Base). These figures are placed here only so the parties can see the order of magnitude the architecture implies and test it against their own thresholds; the true target, attachment point, and glidepath are the first work of accession. The honest purpose of putting numbers on the page is that the viability of the whole instrument turns on whether a rate in this region is one providers will pay and one that funds a layer worth having — and that question cannot be asked, let alone answered, while the number is absent.

9. Assessment Base. In-scope infrastructure and platform cloud revenue, net of resold and intra-group revenue, attributed by customer location. Revenue is the administrable proxy for footprint; the systemic-importance score (clause 6c), not the Base, carries correlation and substitutability into the apportionment. Precise scope is a calibration item (Part K).

10. Perimeter. The Covered Providers are those already designated as critical third parties under the existing regulatory designations of the acceding jurisdictions, aligning the Settlement's perimeter with one States have already drawn.

11. The Fund and its waterfall. The Fund is pre-funded, ring-fenced, and conservatively held, under a published custody mandate and a bar on diversion. Its payout waterfall is: **(i) the Access Floor tranche, ring-fenced as a first call and not available to Coverage; (ii) Coverage, drawn from the balance; (iii) governed surplus.** Because the Access Floor and Coverage draw on the same inflows, the steady-state Assessment (clause 8) must be sized to fund **both** the first-loss layer **and** the committed Access Floor tranche — the two are competing claims on one fund, and the target is set accordingly, not as if either were free. The Fund's accounts are independently audited and published.

Part C — Covered Providers: scoped co-governance and its limits

12. Status. A Covered Provider co-governs the resilience layer — and only the resilience layer. On the technical continuity, redundancy, and portability Standards, where its expertise is real and its interest aligns with every other party's, a Covered Provider holds a genuine seat and vote on the Council's resilience decisions. That co-governance includes a vote on a **resilience-adequacy floor** — a minimum level of Fund resourcing necessary to back the Standards credibly — which a Covered Provider may move to raise. It does **not** include a vote on the **aggregate Fund target** or the Assessment that funds it (clause 13). The two are deliberately distinct quantities: the resilience-adequacy floor is a *minimum the resilience layer may insist upon*; the aggregate target is *the figure that sets the fee*, which may sit at or above that floor and is fixed by the Parties alone. A Covered Provider may argue the floor must be higher; it may not set the rate it pays.

13. The Recusal Bar — what a Covered Provider cannot do, and why. A Covered Provider may not vote on, move, deliberate in, or hold a seat deciding any of the following, and any such vote is void: (a) its own Assessment, or the Assessment rate, base, or methodology generally; (b) the

aggregate Fund target and the glidepath, insofar as these set the aggregate Assessment (the resilience-adequacy *floor* of clause 12 excepted, which it may move to raise but not to bind as the fee); (c) the Continuity-of-Capability Guarantee, Schedule Y, or any service-interruption rule; (d) the expulsion or admission of any Party or Provider; (e) the governance glidepath, voting weights, or this Recusal Bar itself; (f) the Access Floor. **Why the bar is drawn here.** These are the matters on which a provider's interest is structurally adverse to the public interest the Settlement protects — its own bill, the sovereign continuity guarantee, the composition of the body, the public benefit floor. A vote there converts co-governance into self-assessment, the defect that disqualifies the assessed from governing the assessment in every system this instrument is modeled on. The line is not a courtesy that renegotiates upward; it is the boundary that makes co-governance safe to grant at all, and it is itself a Reserved Matter (clause 19). A Covered Provider governs *how the asset is made resilient* and may insist the Fund be *adequate* to that task; it has no vote over *what it pays, whether service may be interrupted, who sits at the table, or what the public receives*.

14. Coverage & Safe Harbor. A paid-up, compliant Covered Provider is a **Covered** entity: (a) **Fund access (first-loss)** — access to the Fund for the defined first-loss layer of a correlated-failure event (clause 8); a mutualized first-loss layer, honestly not full-tail insurance, whose value to a provider lies substantially in discharging or capping first-loss liabilities it would otherwise bear. (b) **Liability safe harbor (accession-conditioned)** — a limitation of liability for covered events that takes effect in an acceding jurisdiction only as and when that jurisdiction enacts it in its own law, on the model of national liability-channeling regimes pairing a cap with a funded compensation route. One jurisdiction's cap does not bind another's courts; the cross-border form is an open problem deferred to Part K, and the safe harbor is worth what each legislature enacts. (c) **Regulatory safe harbor** — coordinated treatment allowing joint resilience action without antitrust jeopardy, within defined limits. The honest statement of the provider's affirmative case to join: co-governance of the resilience layer (Part C); a mutualized first-loss layer (a); regulatory/antitrust comfort for joint resilience action (c) — the most immediate concrete benefit; a liability limitation to the extent each jurisdiction enacts it (b); and, as the genuine macro inducement, a single governed regime in place of an accelerating patchwork of divergent and unpredictable unilateral national regulation. Against that sits a certain, escalating, externally-set cost. Whether the package outweighs the status quo is the providers' decision; these terms set out the case and do not assume its conclusion.

Three structures considered — for the reader's judgment. (A) Seats with full votes — maximizes buy-in and expertise; but lets providers vote on their own Assessment and on sovereign matters, reads as capture, seats private firms as co-sovereigns. Rejected. **(B) Voice without vote** — cleanest optics, no self-dealing; but produces a fiscal hold-up — a provider pays into and is bound by a regime it cannot steer, giving a rational provider reason to refuse and none to want the Fund properly sized. Rejected after the first review identified the hold-up as decisive. **(C) Scoped co-governance with a Recusal Bar** — adopted. Providers co-govern the resilience layer and may insist the Fund be adequate to back it, and are barred from the money and sovereignty questions. This gives a real stake and a reason to want the Fund capitalized, without permitting self-assessment, and reframes the controlling jurisdiction's concession as joining a co-built institution rather than surrendering a chokepoint. **The cost of (C), stated plainly:** it reopens, in contained form, the capture concern (B) foreclosed — a reader may say “providers vote on the standards they must meet, and can push the Fund's adequacy floor up.” The answers are the scope limit (resilience only), the Recusal Bar (clause 13), and the clean split between the *adequacy floor* a provider may raise and the *aggregate target* it may not set (clause 12) — all Reserved Matters. (C) trades a measure of capture-resistance for signability and a coherent motive; the trade is deliberate, and the reader is left to weigh it.

Part D — The Continuity-of-Capability Guarantee

This Part is redesigned from v...b. The prior draft required an Adjudicator to pre-clear a Schedule Y cutoff before it took effect. Adversarial review established that no sovereign will submit its sanctions or national-security actions to prior external clearance — so that design, while strong for the dependent states, was unworkable for the controlling jurisdiction. The protection is therefore moved off the sovereign's act and onto the dependent state's capability: the guarantee is no longer "you may not cut us off without approval" but "a cutoff cannot deprive us of the capability we depend on." This asks far less of the controlling jurisdiction and gives the dependent states a protection that does not depend on an adjudicator being faster than a sovereign.

15. The guarantee. A Party is structurally guaranteed continuity of operational capability notwithstanding any service action taken against it. The guarantee is delivered not by permission but by architecture, through three standing requirements that a controlling jurisdiction cannot unilaterally defeat: (a) **Escrowed continuity** — a Party's essential data and operating state are continuously escrowed or mirrored under independent (non-provider, non-controlling-jurisdiction) custody, so that capability survives a cutoff at the source. (b) **Mandated portability** — Resilience Standards (Part C) require genuine portability and interoperability, so a Party's essential workloads can be moved to and run on alternative capacity rather than being captive to a single provider's switch. (c) **Stand-by capacity** — defined essential-service capacity is maintained in reserve, accessible to a Party on a service interruption, funded as a charge on the Fund or on the providers as part of the Resilience Standards.

16. Schedule Y and what is now permitted. Service interruption remains governed by Schedule Y — the enumerated, exhaustive list of lawful grounds (sanctions compliance, lawful order of a competent court, imminent security threat, non-payment), with a conflict-of-laws rule for competing lawful orders (the rule itself an open problem, Part K). A sovereign may invoke a Schedule Y ground and act on it **without prior clearance** — its sovereignty over its own security and sanctions actions is preserved. What it cannot do is thereby deprive a Party of capability, because the continuity architecture (clause 15) operates regardless of the cutoff's validity. The contest over whether the ground was genuine then proceeds where it belongs — *after* the act, without the targeted Party's lights having gone out in the meantime.

17. Automatic, pre-costed reversal. Any interruption is subject to automatic expedited review by the Adjudicator within a fixed short period. Where the invoked ground is not sustained, reversal is automatic and **remedy is pre-costed and pre-committed** — a defined, escrowed consequence payable to the affected Party without fresh litigation, so that an unjustified interruption carries an immediate and certain price rather than a contingent future judgment. The deterrent is the certainty and speed of the pre-costed remedy, not an adjudicator's prior permission.

18. Non-weaponization and non-discrimination. The Settlement bars use of the infrastructure as a means of coercion among Parties, and requires that the Access Floor, the continuity architecture, and the Resilience Standards apply without discrimination as to a recipient's alignment.

Honest limit (carried, not hidden): continuity-of-capability narrows but does not erase the coercion problem. A determined sovereign can still impose cost and disruption; what it can no longer do is achieve a clean, immediate, total deprivation of a dependent state's capability, because escrow, portability, and stand-by capacity blunt exactly that. The guarantee converts an on/off switch into a friction the targeted state can ride out — which is the protection actually achievable against a party

that physically controls the asset, and is more than prior-clearance could deliver once its emergency carve-out is accounted for.

Part E — Governance

19. The Council and Reserved Matters. Acceding States are voting Parties; Covered Providers vote on the resilience layer only, subject to the Recusal Bar (clause 13). Voting weight among Parties is weighted but bounded: (a) the provider jurisdiction's weight begins as the largest single share and is **capped, with a committed glidepath to no more than 40% by a defined year**, on the model of the satellite consortium that bound allied states into a structure they accepted because dominance was capped and declining; (b) no single Party's weight exceeds the cap regardless of contribution; (c) a **basic-vote floor** gives every Party a minimum weight independent of contribution. To prevent the capped state weight from being reconstituted through its national providers' resilience-layer votes, **provider resilience votes are tallied separately from Party votes and may not be aggregated with the provider jurisdiction's weight**; the resilience layer and the Party governance layer are distinct franchises. **Reserved Matters** — the Assessment rate and base, the aggregate Fund target, the Access Floor, the Continuity-of-Capability Guarantee and Schedule Y, the Recusal Bar, admission and expulsion, and the glidepath — require a supermajority of Parties and are protected from simple-majority change, with a reachable small-Party blocking threshold and no single-Party rewrite.

19A. Accession (the horizon). The settlement is open in principle to any state that accepts and is bound by its core terms — the Assessment, the Continuity-of-Capability Guarantee and Schedule Y, the non-weaponization and non-discrimination rules, and the governance discipline including the Recusal Bar. Accession is conditioned on acceptance of those constraints, not on a state's alignment with the founding cohort. The founding cohort (the provider jurisdiction and its allies) begins the settlement because that is where dependence is deepest and the trust to start already exists; but a closed bloc is not the destination. Admission of a new Party is a Reserved Matter (clause 19), so the founding cohort governs the pace and conditions of widening — but the door is designed open, on the principle that infrastructure the world depends on is more safely governed by a commons that a rival can join on equal terms than by a bloc a rival is built to resist. This is the difference between a defensive pact and a settlement: a rival reduces its own exposure, and the system's, by accepting the same binding terms as everyone else rather than by building a competing stack. That a non-allied state would have reason to accept those terms — less exposure to a chokepoint it does not control, a continuity guarantee, a share in the floor — is the horizon the settlement is ultimately for.

20. The Adjudicator. Disputes under the Continuity Guarantee, post-action review (clause 17), and appeals of Assessment determinations are heard by a neutral Adjudicator. *Identity and seat are an Open Item* (Part I). Because the Adjudicator now acts *after* the fact (not on the critical path of a sovereign act), its independence remains load-bearing but its speed is no longer the single point of failure the prior pre-clearance design made it.

Part F — The Benefit Floor

21. Access Floor. A defined share of Fund inflows is dedicated to a universal digital-access entitlement: affordable basic connectivity and digital public infrastructure in underserved jurisdictions, through a named body, measured against an established affordability metric (entry-level access at no more than 2% of monthly income per capita), without discrimination as to alignment. It is a single

committed entitlement, not a discretionary grant: its share, metric, and non-discrimination rule are Reserved Matters, and its tranche is **ring-fenced as a first call on the Fund, segregated from and ahead of Coverage** (clause 11), so a Coverage event cannot defund it. Because floor and Coverage compete for one inflow, the Fund target (clause 8) is sized to carry both; the floor is protected first, and Coverage is sized on the balance.

Part G — Durability & Commitment

22. Durability undertakings. The Settlement's value to every counterpart depends on the provider jurisdiction sustaining its commitments across changes of government. To make durability auditable rather than asserted, the provider jurisdiction evidences it through an observable package: (a) the Assessment and Fund established in statute, not executive act alone; (b) the Fund ring-fenced in law against diversion; (c) the Settlement entered by the most durable available domestic ratification; (d) a defined snap-back consequence for failure to maintain the undertakings. These speak to **durability** — whether the commitment holds against drift, neglect, or quiet defunding — and are honest about their limit: no civil undertaking deters a sovereign that has decided the geopolitical stakes exceed the penalties. Durability is guarded here; **capability against misuse** is guarded by Part D. The two are carried separately, because the domestic majority that supplies durability does nothing for non-abuse, and a counterpart should verify each independently.

Part H — What Each Party Receives

Party
Gives
Receives

Provider jurisdiction (a willing future government — Recital F)

Caps and dilutes its governance weight; binds itself to the continuity architecture and the Durability undertakings — but **retains its sovereign right to act on Schedule Y grounds without prior clearance**

Membership in a co-built resilience institution in place of eroding sole control; allied cohesion that slows defection to sovereign alternatives; a priced settlement of a risk now carried ungoverned

Acceding States

Assessment buy-in; acceptance of weighted governance

Continuity of capability that survives a cutoff *by architecture* rather than by an adjudicator's speed; portability and stand-by capacity they cannot buy privately; a governance seat with a protected floor; a share in the Access Floor

Covered Providers

The Assessment; acceptance of Resilience Standards; the Recusal Bar

Co-governance of the resilience layer and a vote on its adequacy floor; a mutualized first-loss layer; regulatory/antitrust comfort; a liability limitation to the extent enacted; a single governed regime in place of a divergent unilateral patchwork

Underserved jurisdictions / the public

—
A funded, ring-fenced, non-discriminatory universal-access floor with a defined metric and named deliverer

Part I — Open for Negotiation

the exact entry rate within the 1–5 bps band; the Fund target, first-loss attachment point, and glidepath (the illustrative clause-8 figures are for testing only);
the individual-provider weightings within the systemic-importance score;
the precise Assessment Base scope;
the contents of Schedule Y and the conflict-of-laws rule;
the engineering and custody of the continuity architecture (escrow, portability standard, stand-by capacity) and the pre-costed remedy schedule and review clock (Part D);
the resilience-adequacy-floor mechanism and its relation to the aggregate target;
starting governance weights, the cap year, the basic-vote floor, Reserved-Matters thresholds, and the separate-franchise tally;
the Access Floor share and named delivery body;
the identity and seat of the Adjudicator;
accession terms and sequence.

Part J — Basis & Precedent

The Assessment as systemic-risk fee — deposit-insurance assessments supply the *form*; a pre-funded resolution fund supplies the *absolute rate* (target-sets-rate); the systemically-important-bank ranking supplies only the *apportionment among providers*. The bank surcharge is a capital requirement, not a fee, so only its ranking logic transfers, and only to distribute the Assessment. Risk-based deposit insurance already operates exactly this way — a pool target sets the aggregate, a risk score sets each payer's slice.

Coverage as a first-loss layer — post-event public backstops re-opened uninsurable perils by absorbing and pricing a defined layer, letting a private market re-form above it. (Honest note: the classic backstop takes the *tail* and leaves first loss to the market; this Fund takes a *first-loss* layer to seed a market for a peril that has none. The inversion is deliberate and is the contestable design choice a reviewer should press.)

Continuity-of-capability — escrow, portability, and stand-by capacity as resilience guarantees draw on data-portability and operational-resilience regimes and on the logic of source-independent continuity; the protection is structural, not permission-based.

Capped, declining dominance with a basic-vote floor — the satellite consortium's capped glidepath and the weighted-voting institutions' basic vote; both held where the asset was *co-owned*, which scoped co-governance (Part C) is what supplies here.

The benefit floor — the universal-service tradition, the broadband-affordability metric, and the digital-public-infrastructure model proven at national scale.

The backing analysis carries sources, figures, and provable/derived/inference labels.

Part K — Deliberately Deferred to Detailed Work

Their absence here is deliberate; two are genuine open problems, marked: - provider-by-provider systemic-importance scoring and Assessment apportionment; - the Fund target, first-loss attachment point, and stress-tested sizing (the clause-8 figures are illustrative, not these); - **(open problem) the cross-border legal form of the liability safe harbor** — channeling and a compensation route enacted across multiple sovereign systems; - **(open problem) the conflict-of-laws rule in Schedule Y** — competing lawful orders across jurisdictions; clause 16 depends on it; - the engineering, custody, and cost allocation of the continuity architecture, and the pre-costed remedy

schedule and review clock (Part D); - the legal form of accession in each jurisdiction. This document establishes that the architecture is coherent, precedented, and worth that work — not that the work is done.

Note on status, method, and version

Status. Indicative and pre-decisional. Non-binding. Circulated in confidence for reaction only. No clause above commits any party; figures are entry parameters and illustrative test values for discussion, not demands.

Method. Prepared under a receipts-first standard, with AI research tools used as a function of time and scope. Load-bearing judgments are labeled provable, derived, or inference in the backing analysis; the weak points are named rather than smoothed; and the structure was put through three rounds of adversarial review, with what broke cut and recorded. The aim was work that can be checked rather than trusted. The backing analysis is available on request.

Version. cloud-grid-termsheet_v2026-06-22c. This version, in response to the third review, replaced the continuity guarantee's pre-clearance mechanism with a continuity-of-capability circuit-breaker plus automatic pre-costed reversal (Part D, because external pre-clearance of a sovereign security act is not signable for the controlling jurisdiction); placed an indicative Fund target and first-loss attachment point on the page as illustrative test values (clause 8); split the Fund-target recusal into a votable resilience-adequacy floor and a recused aggregate target (clauses 12–13); and reconciled the Access Floor and first-loss layer as competing draws on one inflow (clauses 11, 21).