ECONOMIC VIABILITY AUDIT

MintID (Agent-KYC / KYA)

Independent Economic Viability Audit

Opinion: Qualified — see §2.1

ENGAGEMENT Independent economic, tokenomic and survivability audit of the MintID concept and its implementation plans

AUDITOR Conectia PRO

SUBJECT CORPUS master/ (frozen concept and specification set) plus supporting archive documents

AUDIT PERIOD July 2026

DELIVERED 2026-07-30

REPORT LANGUAGE English (source corpus is in English; all citations refer to the original documents)

STATUS Final — suitable for attachment as a product-viability audit annex

Audit performed by Conectia PRO · July 2026 · delivered 2026-07-30 · conectia.pro

1. Scope, methodology and limitations

1.1 Engagement scope

This audit examines the **economic viability, tokenomics and survivability** of the MintID concept — a sovereign Layer-1 blockchain (Cosmos SDK / CometBFT) providing privacy-preserving, KYC-backed identity for AI agents, with self-custodied USDC/EURC escrow, a dual-rail disclosure model (commercial / dispute), an arbiter market, and a native token with a ratified hard cap of 108,000,000 units.

The audit adopts the perspective of **a prospective investor and a prospective validator ("miner")** asking: *Is this economically viable? Can each participant in the system be solvent? Can the project survive to and beyond mainnet?* It concludes with the **claims and corrections the auditor considers necessary** to secure the project's continuity and the economic solvency of its stakeholders.

1.2 Documents audited

Document	Version / freeze	Role
KYA-Concept-and-Market-Paper.md	v1.3, frozen 2026-07-14	Concept & market whitepaper (non-normative)
Agent-KYC-Roadmap.md	Jul 2026 (updates to 2026-07-20)	Committed launch scope, hard launch gates
SPEC-1-Sovereign-Privacy-Preserving-Identity-Blockchain.md	v1.3, frozen 2026-07-20	Normative chain spec: consensus, token, validators
SPEC-2-Agent-KYC-and-Verifiable-Human-Backed-Identity.md	current frozen	Normative agent/escrow layer
SPEC-3-Consensual-Disclosure-and-Graduated-Remediation.md	current frozen	Normative disclosure/dispute/arbiter layer
Genesis-Allocation-and-Transparency-Policy.md	v1.1, 2026-07-14	Genesis allocation, MiCA posture, entities
Use-Case-x402-Agent-Payments.md	MINT-144, 2026-07-13	Flagship market use case
Agent-KYC-Phase-0-Backlog-Plan.md	current	Implementation plan (~90 cards)
archive/Agent-KYC-Concept-and-Market-Paper-v2.md	archived	Only extant text of R-A13/R-A14/R-A15
archive/Agent-KYC-Model-Weak-Points-Review-v5.md	archived 2026-07-13	Risk register (referenced)

1.3 Methodology

1. **Full documentary review** of the corpus above, with line-level citation of every material claim.
2. **Internal-consistency analysis**: cross-document reconciliation of every economic mechanism, figure, requirement ID and gate.
3. **External market validation** (web research, accessed July 2026) of every market figure the corpus relies upon, plus reference pricing for the cost items the corpus leaves unbudgeted (audits, KYC checks).
4. **Derived quantitative analysis** where the corpus provides inputs but no computation (unit economics, cost incidence). All derivations are labelled as auditor arithmetic.
5. **Stakeholder solvency mapping**: cash-flow in/out and risk exposure for each of the seven economic roles.

Throughout this report, three epistemic categories are kept distinct: **[D]** documented fact (cited to file:section/line), **[A]** auditor arithmetic derived from documented inputs, and **[J]** auditor judgment.

1.4 Limitations

- This is **not** a code audit, a cryptographic audit, or a legal opinion. Where the corpus itself gates items on such reviews (C2-P0, NC2, counsel items), this report treats those gates as costs and dependencies, not as subjects.
- The repository's implementation directories were not assessed for code quality; the Phase-0 backlog itself records the starting point ("scaffolding only") and is audited as a plan.
- External figures were validated against publicly available secondary sources during July 2026; primary paywalled analyst reports (IDC, Omdia, Gartner originals) were not purchased.
- The corpus is explicitly a **pre-genesis, pre-parameter** design ("mechanism committed, numbers gated"). This audit therefore evaluates (a) whether the *mechanisms* are economically sound, and (b) whether the *process that will produce the numbers* is adequate, funded and sequenced. Where numbers do not exist, the audit says so rather than inventing them.

---

2. Executive summary and audit opinion

2.1 Opinion

Qualified opinion. The MintID economic design is unusually disciplined for a pre-genesis crypto project: it refuses to publish invented numbers, hard-caps supply, bans price-support and discretionary post-genesis minting, gates every parameter on an independent economic simulation, and subordinates its own genesis allocation to a third-party economic audit (R20) *before genesis exists*. Several mechanisms — the burn-only rule for unclaimed remediation, native 4-year vesting, the no-listing MiCA posture, the no-privileges first-implementer stance — are best-practice and should be preserved.

However, in the auditor's opinion the project is **not currently financeable or launchable as documented**, for five critical reasons:

1. The genesis allocation policy contains a **numerical contradiction of a factor of 11–23×** (a declared 10–20% development band versus a 0.93% working figure described as "within the band") on a decision that is **irreversible by construction** (no post-genesis mint) — Finding **C-01**.
2. The launch plan rests on a **circular dependency the documents themselves create**: no listing → no market price → the §5.4 minimum-economic-security launch gate ("bonded stake *value*") is unmeasurable → the issuer bond has no known deterrent value → professional validators cannot be recruited under a "never-promised, retroactive-only" reward rule — Finding **C-02**.
3. The **entire funded-accountability layer hangs on requirements (R-A13/R-A14/R-A15) that have no normative text in any live document**; the only extant text lives in a superseded archive draft that mandates the abandoned fiat-custodian model — Finding **C-03**.
4. **No economic parameter exists** beyond MAX_SUPPLY: no fee, band, bond size, protocol cut, emission curve or burn split — and the independent simulation that must produce all of them, itself on the critical path to genesis, **had not been commissioned** as of the corpus freeze — Finding **C-04**.
5. **Protocol revenue is structurally decoupled from adoption**: verification is free by construction, state and fees scale with the small bounded set O(issuers + verifiers + validators), the only scalable revenue line (premium services, R26) is barred from accruing value to the token, and the one population-scaling lever (agent mint/lease) is non-normative and parameterless. When the emission pool is exhausted, the validator security budget has no documented successor — Finding **C-05**.

None of these five is unfixable, and none invalidates the underlying market thesis — which external validation broadly **supports** (§4). But all five must be remediated **before** genesis parameters freeze, and three of them (C-01, C-02, C-03) are pure documentation/design decisions that could be closed in weeks at near-zero cost. The conditions-precedent matrix in §10 defines what a favourable opinion would require.

2.2 Top findings at a glance

ID	Severity	Finding (short)
C-01	Critical	Genesis allocation band contradiction (10–20% vs 0.93%), irreversible decision
C-02	Critical	Circular launch dependency: no market price → unmeasurable security gate → unrecruitable validators
C-03	Critical	R-A13/R-A14/R-A15 have no live normative text; archive text contradicts current design
C-04	Critical	Zero ratified economic parameters; critical-path simulation not commissioned
C-05	Critical	Protocol revenue structurally bounded; no token value accrual from scalable lines; post-emission security budget undefined
H-01	High	Adverse unit economics against the flagship x402 market (~$0.30/tx, ~$725 lifetime volume/agent)
H-02	High	Arbiter market cold-start: launch-blocking role paid only by dispute losers, zero volume at launch
H-03	High	Large unbudgeted cost stack (audits, counsel, bounty, simulation, foundation) with no documented fiat funding source
H-04	High	Genesis accounts can delegate and vote: §5.4 thresholds satisfiable with insider stake
H-05	High	Circular calibration dependencies: fees ⇄ simulation ⇄ pilot ⇄ live product ⇄ genesis
H-06	High	Issuer cost structure forces concentration at the identity-link holder; first-implementer related-party conflicts undisclosed as such
H-07	High	Competitive timing: transparent-identity rails (KYA/Skyfire, ERC-8004, Visa/Mastercard, x402) shipping now — externally confirmed
M-01…M-10	Medium	See findings register (§9)

2.3 What the external validation showed (summary)

- The demand-side market thesis is **corroborated**: IDC's >1B agents by 2029, Gartner's 40%-of-enterprise-apps by 2026, Omdia's 50:1 NHI ratio and the x402 figures all check out against independent sources (§4).

- The **competitive timing risk is confirmed and has worsened**: by April 2026 "every major payment network had shipped or unveiled a KYA primitive"; Skyfire has shipped Visa Intelligent Commerce integration and an F5 enterprise partnership.
- The **KYC per-check price anchor** the whitepaper leaves as an open task is unfavourable at the low end: digital IDV averages ≈ $0.20/check (Juniper), with full KYC bundles $1.50–$4.00 — which compresses the room between "mint fee low enough for agent unit economics" and "mint fee high enough to be a sybil floor" (§6.2).
- The **agent-insurance market figure** carried in SPEC-3's deferred extension ($4.48B 2025 → $25B 2030) is **not corroborated**; independent estimates for 2025 range $0.3B–$1.2B. This affects only the deferred insurance extension, but the figure should not be cited to partners.

---

3. Project and business model overview

3.1 The concept in one paragraph

A human (or organisation) completes KYC **once** with an issuer; the resulting privacy-preserving credential backs **every AI agent** that principal deploys. An agent can prove, in zero knowledge, *"I am operated by a KYC-verified, liable human, at assurance grade ≥ G, within scope S, with ≥ F of pre-funded recourse behind me"* — without revealing who the human is, and without the principal's agents being linkable to each other. Disputes settle **money-before-identity**: escrow pays out with no identity revealed (L0); a minimal fraud signal escalates without PII (L1); full unmasking happens only under due process (L2). Separately, a **consented commercial rail** lets a relying party *pay the principal* for graduated identity disclosure the principal may always refuse. (KYA-Concept-and-Market-Paper.md §0, §3.3–3.4.)

3.2 The five committed revenue flows (§8.1 of the whitepaper)

Flow	Mechanic	Who earns	Token linkage
A · Agent mint + lease	Per-agent mint fee + recurring lease	Issuer (+ protocol cut, partially burned per R17)	Only via undefined "protocol cut"
B · Issuer economics	Registry admission/renewal + slashable bond (R16)	Protocol (fees partially burned); issuer charges its own clients	Direct but small-N
C · Disclosure brokerage	disclosure_bid paid to the principal; issuer brokerage fee; optional protocol cut	Principal + issuer	Only via undefined cut
D · Dispute / arbiter market	Bonds, arbitration fees borne by the loser	Arbiters (off-chain service revenue)	None by rule (SPEC-1 §6.1)
E · Premium verification services	Mirrors, archives, SLAs, paid proof APIs (R26)	Operators	**None by rule** ("must not influence consensus rewards")

[J] The flow architecture is coherent as a *services* economy, but as a *token* economy it has a designed-in gap analysed in §5.3: flows D and E are firewalled away from the token by explicit rule; flows A and C touch the token only through a protocol cut whose percentage does not exist anywhere in the corpus.

3.3 The amortisation argument

The whitepaper's core unit-economics claim (KYA-Concept-and-Market-Paper.md §8.1): one human KYC (market price: see §4, $0.20–$4.00 per check; full onboarding higher) amortises across every agent the human deploys, so the marginal cost of *accountability per agent* tends to the mint cost, not the KYC cost. With NHI:human ratios of 50:1 today and agent populations forecast in the billions, per-check models re-pay verification per agent per counterparty while MintID pays once per human.

[J] The argument is directionally sound and is the strongest economic idea in the corpus. Two caveats materially weaken it as stated: (a) the low market price of a *basic* IDV check means the amortisation saving per agent is cents-to-dollars unless the comparison anchor is full compliance onboarding (which the whitepaper acknowledges as an unverified open task — "verify per-check market prices against primary sources before publication"); (b) the saving accrues to the *principal*, while the protocol monetises only the mint/lease — so amortisation is an adoption argument, not a protocol-revenue argument.

3.4 Competitive position

The corpus's own competitive register (§2.4, §13.3) is honest and current: KYA/KYAPay (Skyfire) ships today with transparent bearer JWTs; Concordium runs a live identity-L1 with an agent registry; ERC-8004 gives the transparent camp an open standard; Visa TAP and Mastercard Agent Pay own the card-adjacent segment; x402 is the fastest-growing agent-payment rail with **no identity layer at all** — which is precisely MintID's product statement. MintID's differentiation — privacy-by-default, funded recourse, metered identification with no tracking rights — is real and is not served by any incumbent. The risk is not differentiation but **timing** (Finding H-07).

---

4. Market validation (external, accessed July 2026)

Every market figure the corpus relies on was checked against independent public sources. Verdicts: **Corroborated** (independent source matches), **Partially corroborated** (directionally supported, magnitude uncertain), **Not corroborated** (independent sources diverge materially).

#	Corpus claim	Corpus source	External check (July 2026)	Verdict
1	>1B active AI agents by 2029 (from ~28.6M in 2025)	IDC (Rick Villars, Dec 2025), cited whitepaper §8.1/ §13.3	IDC forecasts >1B deployed agents by 2029 (40× 2025); related IDC projections reach 2.216B by 2030, 217B daily actions	**Corroborated**
2	40% of enterprise apps with task-specific agents by 2026, from <5% in 2025	Gartner, cited §8.1	Gartner press release (2025-08-26) states exactly this	**Corroborated**
3	NHI:human ratio 50:1	Omdia Dec 2024 via IBM Think (2026-02-02)	Omdia 50:1 confirmed via IBM and Dark Reading; other estimates range 40:1–82:1	**Corroborated** (band is wide)
4	x402: ~69k active agents, ~165M transactions, ~$50M cumulative volume (late Apr 2026)	Use-Case-x402 §2.3, flagged "dated claims"	Coinbase-reported figures match; Chainalysis reports 100M+ agentic payments on Base; ~$600M annualised volume by Mar 2026; **average transaction $0.20–$0.30**	**Corroborated**; avg-tx figure independently confirms the auditor's $0.30/tx derivation

#	Corpus claim	Corpus source	External check (July 2026)	Verdict
5	KYC/IDV per-check market price (whitepaper's open "comparison anchor" task)	Not yet sourced in corpus	Juniper: ≈$0.20 avg per digital IDV check (2025), falling to ~$0.17 by 2029; document checks $0.10–$1.50; biometric $0.25–$2.00; bundled KYC $1.50–$4.00/ approved user; vendor minimums $20k–$100k/yr	**Now sourced** — see §6.2 for the adverse implication
6	Kleros lineage viability: "~3,000 cases resolved"	SPEC-3 :415	No current public total confirmed; Kleros active (2026 project updates, municipal ADR pilots) but case-volume figure unverifiable from public secondary sources	**Partially corroborated** — treat count as unverified
7	Agent liability insurance market $4.48B (2025) → $25B (2030)	SPEC-3 :414 (deferred extension only)	Independent 2025 estimates: $0.3B (Fact.MR), $1.2B (MarketIntelo); agentic-AI-insurance (broader) $18.16B by 2030	**Not corroborated** — divergent by 4–15×; do not cite externally
8	KYA/Skyfire ships today; timing risk	Whitepaper §2.4	Confirmed and stronger: F5 partnership, Visa Intelligent Commerce demo (Dec 2025), ERC-8004 on-chain coupling, "by April 2026 every major payment network had shipped or unveiled a KYA primitive"	**Corroborated — risk worse than stated**
9	Rented security (ICS/ Babylon) is a viable option for young chains	SPEC-1 §5.4.3	Babylon/ICS actively marketed to exactly this profile ("early-stage networks face a security problem before a growth problem"); no public standard pricing	**Corroborated qualitatively**; cost remains unbudgeted (M-01)

#	Corpus claim	Corpus source	External check (July 2026)	Verdict
10	ZK-circuit and contract audit market costs (unbudgeted in corpus)	Not in corpus	Smart-contract audits $5k–$250k; DeFi protocols typically $25k–$100k; **ZK/ novel-crypto audits routinely >$150k** with ~3× complexity premium; remediation passes $5k–$20k each	**Now sourced** — feeds the unbudgeted-cost stack (H-03)

Sources: IDC via InfotechLead, IDC blog, Gartner press release, IBM Think NHI guide, Dark Reading / Omdia, Chainalysis x402, RZLT x402 explainer, ComplyCube KYC pricing, Didit pricing, fintechspecs KYC benchmarks, QuillAudits pricing, Sherlock audit pricing, Kleros blog, Fact.MR, MarketIntelo, Skyfire, Skyfire × Visa, F5 × Skyfire, Eco KYA explainer, Babylon, Cosmos forum.

Assessment [J]: The demand thesis survives external scrutiny — the market MintID targets is real, large and growing on the corpus's own cited trajectory. What external validation *adds* against the corpus is (i) confirmation that the transparent-identity competition is consolidating faster than the whitepaper's frozen text records, (ii) a now-sourced KYC price anchor that is lower than the amortisation narrative implicitly assumes at its bottom end, and (iii) hard reference costs for the audit stack the plans leave unbudgeted.

---

5. Tokenomics analysis

5.1 Supply, emission and burns — what is fixed and what is not

Fixed [D]: - MAX_SUPPLY = 108,000,000, immutable hard cap (SPEC-1 §6.1 L315), ratified 2026-07-14 (MINT-150), superseding 93,924,792 ratified two days earlier (MINT-118). No economic rationale for either figure, or for the ~15% change between them, appears anywhere in the corpus. - REWARD_POOL = MAX_SUPPLY – GENESIS_MINTED; emission E(epoch) = min(remaining_pool, scheduled_declining_epoch_emission) (§6.1 L319). - Burns never refill the reward pool (L328); realised emission can only undershoot schedule (L329); inflation defaults zeroed at genesis (L331); explicit refusal to promise perpetual deflation (L36) and to tax validator stake to force deflation (L346).

Not fixed [D]: decay factor, epoch length, initial emission, horizon, GENESIS_MINTED, every fee, every burn split ("defined portion", undefined), the protocol rent, the token denomination and exponent. All delegated to the Phase-0 independent simulation (SPEC-1 L324, L1252; Open policy parameters #1).

Assessment [J]: The monetary *rules* are conservative and well-drafted; the refusal to promise deflation is honest. But an investor is being asked to underwrite a monetary policy of which only the asymptote (108M) exists. The cap itself changed by 15% in two days without recorded justification — harmless now, but it demonstrates that the number is not anchored to any model yet. Until the simulation exists, the tokenomics are a **shape, not a policy**.

5.2 Validator security budget — the survivability core

[D] Validator rewards come exclusively from emission: R_i = E(epoch) × W_i / ΣW (§5.2 L266–283). No documented share of transaction fees flows to validators. Off-chain service revenue (verification APIs, arbiter fees) "must never influence consensus rewards" (R26 L77, SPEC-1 §6.1). The active set targets 100 validators (§5.1 L256).

[A] Consequence chain: the security budget is a **strictly declining** function of time by construction (declining emission, hard cap, burns that never re-mint). At pool exhaustion the budget is whatever a fee market provides — and the fee market is levied on

the bounded set O(issuers + verifiers + validators + roots), *not* on usage (verifications are free, presentations never touch the chain — R15/R18, §10.5 L806). With a plausible early population of tens of issuers and a governed lease schedule that does not yet exist, there is no documented path to a nine-figure-token-cap chain paying 100 professional validators from fees.

[D] The corpus knows the young-chain problem and answers it procedurally: the §5.4 minimum-economic-security launch gate (bonded stake value, independent operator count, Nakamoto floor — thresholds from the simulation, same rank as R20), the validator-bootstrap plan, the rented-security evaluation (ICS/Babylon, Must-evaluate before genesis), and early-epoch emission weighting toward validators (§5.4.4 L306).

[J] The §5.4 apparatus is genuinely good process design — but it currently rests on two unresolved defects: the gate is denominated in a **market value that cannot exist under the project's own no-listing policy** (Finding C-02), and the recruitment plan's incentive mechanism is **legally required to promise nothing** (Genesis Policy §1.6: "retroactive, discretionary grants only — never promised token rewards"). A professional Cosmos operator cannot budget genesis infrastructure against an unpromised, discretionary grant of an unpriced asset. As written, the plan asks validators to be patrons, not counterparties. Additionally, §5.4.4's instruction to weight early emission "toward validator rewards rather than other uses of the reward pool" references reward-pool uses that no other section defines (M-02).

5.3 Token value accrual — the designed-in gap

[D] The token's stated role: "the economic security asset… validator bonding, delegation, issuer bonds, transaction fees, state/registry leases, governance-weighted economic operations" (§6.1 L312). The scalable economic activity of the network, however, is deliberately routed **off-chain and off-token**: - Verification: free by construction (R15/R18) — zero protocol revenue, zero burn, on the network's dominant activity. - Premium services (R26): direct operator fees, barred from consensus rewards, no protocol cut, no burn. - Arbiter fees: off-chain service fees, same bar (SPEC-3 R-A28, SPEC-1 §6.1). - Disclosure bids: paid to the *principal*; issuer takes brokerage; only an **optional** protocol cut, "partially burned", percentage undefined (whitepaper §8, SPEC-3 §8). - Agent mint/lease: the one flow that scales with population — non-normative in SPEC-1 (L472), parameterless, with its launch-blocking disclosure-policy item still open (L1263).

[A] Result: of the five flows in §3.2, the two that scale with adoption (D, E) are firewalled from the token by rule; the two with protocol cuts (A, C) have undefined percentages; the one fully on-chain flow (B) is proportional to the number of *issuers* — a set the design keeps small on purpose. Token demand therefore rests on: validator/delegator bonding, issuer bonds (min_bond placeholder ≈ 1000 tokens ≈ 0.000926% of supply), gas on a low-transaction chain, and governance.

[J] This is a **structural value-accrual gap, not an oversight** — the firewall exists for a good reason (keeping consensus neutral, SPEC-1 §6.1). But the consequence must be priced: the token is a pure security-and-access asset on a chain whose fee base is bounded by design. For an investor, the honest framing is that flows A and C's protocol cuts and burn percentages are **the entire token investment case**, and they are currently blank. The remediation is not to break the firewall but to make the simulation deliver a **protocol P&L** (issuers × fees, agents × mint/lease, disclosure volume × cut, burn totals vs emission) under published adoption scenarios, and to ratify non-zero protocol-cut bands for flows A and C before genesis (§9, C-05; §10, EC-1).

5.4 Genesis allocation and insider-stake integrity

[D] Genesis policy: single genesis event, no post-genesis discretionary mint (§2.2 L31); development allocation finalised **as a percentage of MAX_SUPPLY within the market-conventional 10–20% band** by the simulation (§2.1 L30); working figure "on the order of 1,000,000 tokens", with the policy's own illustration: "1M ≈ 0.93% — conservative within the band" (§2 L28, §2.1 L30); 4-year/1-year-cliff native vesting (§4.2 L58); 12-month lock-up beyond vesting (§4.4 L66); labelled addresses, hash-committed allocation statement (§4.1 L54); R20 economic audit covers the allocation before genesis (§4.5 L70).

Two defects:

1. **C-01 [D/A]:** 0.93% is not within a 10–20% band; the two anchors differ by 11–23×. Because §2.2 forbids any later mint, whichever reading wins is permanent. At 1M tokens the development entity is plausibly under-capitalised for a multi-year, audit-heavy, multi-entity programme (see H-03); at 10–20% the "conservative" narrative and every published communication based on the 1M figure is wrong. This must be resolved as an explicit, justified decision — not left for the simulation to discover.

2. **H-04 [D/A]:** vesting and lock-up restrict **transfers**, but nothing in either document restricts **delegation or governance voting** by genesis accounts — and Cosmos SDK vesting accounts can delegate. The §5.4 gate's thresholds (bonded-stake value, operator count via delegations, Nakamoto floor) can therefore be satisfied in part by the foundation delegating to itself or to affiliated validators. The policy half-recognises this (§2.1 asks the simulation to model the allocation's "impact on early staking security") but imposes no rule. Required: an explicit no-delegation/no-vote commitment for genesis accounts, or their exclusion from gate arithmetic (§10, EC-4).

5.5 MiCA posture and its economic side-effects

[D] The compliance design is prudent and coherent: testnet tokens worthless by construction; genesis allocation under service agreements is not an offer to the public; the reward-pool emission rides the validation-rewards exemption; **no listing** until the full Title-II path is deliberately executed (CNMV notification in Spain); no voluntary "MiCA-style" white paper (Genesis Policy §3 L36–48).

[J] The economic side-effect is the C-02 circularity: "no dump by construction" also means **no price discovery by construction**, which unprices the launch gate, the issuer bond, validator compensation and the token component of any contributor package (Spanish tax on vested tokens is flagged (counsel) but valuation-at-vesting of an unlisted token is also an unsolved economic input). The audit does not recommend abandoning the no-listing posture; it recommends **denominating every launch-critical threshold in measurable units** (fiat-denominated attack cost via a published proxy, operator counts, Nakamoto coefficient, bonded percentage of supply) rather than market value (§10, EC-3).

---

6. Unit economics and capital costs

6.1 Per agent — tested against the corpus's own flagship market (x402)

[D] x402 late-April 2026 figures carried in the corpus: ~69,000 active agents, ~165M transactions, ~$50M cumulative volume; externally corroborated, with independent confirmation that the **average transaction is $0.20–$0.30**.

[A] Derived: ≈ $0.30/transaction; ≈ 2,391 transactions/agent; ≈ **$725 cumulative volume per agent**. Against that revenue base, a MintID-verified agent bears: mint fee (undefined) + recurring lease (undefined) + escrow lock gas + a CCTP funding cycle of ~19–21 minutes per funding event (SPEC-2 :367) + **capital F immobilised for the 3–12-month credential life at zero yield** (per-agent pre-allocated funding, the pooled alternative being deferred; float yield removed with the fiat-custody pivot). Even at a nominal F of a few hundred dollars, the opportunity cost of locked capital is of the same order as the agent's entire gross transaction volume.

[D] The corpus itself registers this as an unreviewed gap: stopper **S-8** / decision **D-x5** ("fee economics at x402 fleet scale — mint/lease vs sybil floor") assigned to the same uncommissioned simulation.

[J] Finding H-01: for the micro-transaction segment the corpus presents as its fastest-growing wedge, the committed capital model is presumptively uneconomic below a per-agent value threshold that nobody has computed. This does not kill the concept — high-value segments (B2B agents, escrow-backed commerce, regulated flows) plausibly clear the bar — but the **x402 identity-MVP go/no-go (MA3) should be gated on a published unit-economics study**, not discovered in production (§10, EC-7).

6.2 The sybil floor vs the price ceiling

[D] Anti-sybil rests on: valid human KYC + per-principal active-agent cap + KYC freshness + "mint fee/lease economics" (R-A7), all enforced privately per issuer; **cross-issuer evasion is explicitly unmitigated research** (SPEC-2 :82, :349). The whitepaper's comparison anchor — the market price of a KYC check — is now sourced: ≈ $0.20 average per digital IDV check, $1.50–$4.00 bundled.

[A] The effective sybil cost is min(KYC cost + mint fee) *across issuers*. The mint fee is thus squeezed from both sides: it must be **low** relative to a $0.30/tx, $725-lifetime-volume agent economy (§6.1), and **high** enough to price out sybil fleets whose

alternative cost is a $0.20–$4 KYC check per fake principal. The corridor between those constraints may be empty for low-value segments; it has never been computed.

6.3 Per issuer — fixed costs force concentration

[D/A] An issuer must: publish a status root every 30 seconds (R9) — ≈ **2,880 chain transactions/day in perpetuity**, gas economics of the permissionless relayer unspecified; run the emergency-root path; operate a revocation push channel that carries "zero security weight" (pure cost, R-A6.1); run AML provenance checks per mint (Chainalysis/TRM-class per-query fees); custody the identity-link vault (R-A3); operate the disclosure broker and consent engine; and post a slashable bond whose deterrent value is unpriced (C-02). These costs are **fixed per issuer and independent of portfolio size**, while revenue (mint, lease, brokerage — all price-undefined) scales with portfolio.

[J] Finding H-06: strong economies of scale at the exact node holding the only identity link → structural pressure toward few, large issuers — in tension with the multi-issuer neutrality argument used against KYA (§13.1 of the whitepaper), and compounded by the fact that the first issuer, the reference SDK/service operator (R26 revenue) and the development company are the **same corporate group** (CONECTIA OÜ / Spanish SL), with the foundation as steward. The corpus handles the *privilege* question well (same bond, same admission, verifiable on-chain — §14.3); it does not yet handle the *related-party economics* question (transfer pricing is flagged (counsel); revenue-concentration disclosure is not). An investor-grade information pack must disclose these related-party flows as such (§10, EC-10).

6.4 Per principal — capital immobilisation and dispute exposure

[D] Per-agent pre-allocated funding: capital = $\Sigma F_i \times$ credential life (3–12 months), zero yield, per agent (SPEC-2 :319, :360). Loss of sk_P under the recommended k_release construction leaves discretionary withdrawal blocked until credential expiry (funds delayed, not lost — :432). A contested claim **freezes the reserve for the whole dispute including the judicial phase**, survives credential expiry and self-revocation (SPEC-3 :180, :200), and the escalation ladder's accumulating-cost design means the loser pays everything — which is also "cheap to be rich": a deep-pocketed counterparty can rationally push a marginal case up the ladder against a capital-constrained principal (M-06).

6.5 Per arbiter and per verifier

[D] Arbiters: registered, KYC'd, bonded, stake-sortitioned, slashable — and **paid exclusively from fees borne by dispute losers** (SPEC-3 R-A28). Launch-blocking for the F2+ dispute path. **[A]** At launch, dispute volume ≈ 0 → arbiter income ≈ 0 → no rational actor posts a bond and maintains KYC'd readiness → the launch-blocking role has no supply side. No subsidy, retainer or bootstrap budget exists anywhere in the corpus (Finding H-02).

[D] Verifiers: two tiers — bonded (permanent slashable bond, standing L0-claim access) and permissionless (no bond, JIT claim bond sized to the claim, slashed only if frivolous) — verifying identically on the 10-second path (SPEC-2 R-A11.1). **[J]** The permissionless tier offers the same dispute access at strictly lower carrying cost; unless bonded status confers priced advantages (rate limits, claim latency, seller-trust signalling), the bonded tier is economically dominated and its deterrence capital never accumulates (M-07).

---

7. Stakeholder solvency analysis

For each economic role: inflows, outflows/risks, and the **solvency condition** this audit considers necessary. "Undefined" means no figure exists in the corpus.

7.1 Investor (token holder / backer of the operating group)

- **Inflows:** token appreciation via undefined protocol cuts on flows A/C plus burn pressure (undefined splits); equity-side revenue from R26 services and issuer operations (accrues to CONECTIA group, not the token).

- **Outflows / risks:** funds the entire pre-genesis cost stack (H-03); no liquidity by construction (no listing); cap changed once without rationale; genesis allocation ambiguous by 11–23× (C-01).
- **Solvency condition:** a ratified protocol P&L with non-zero cut/burn bands (EC-1), resolution of C-01 (EC-2), and a disclosed fiat budget/runway (EC-5). Until then the token is unpriceable and the investment case is entirely equity-side.

7.2 Validator

- **Inflows:** declining emission share only; no fee share documented; possible retroactive discretionary grants (unpromisable by policy).
- **Outflows / risks:** genesis infrastructure at own cost; slashing; an asset with no market and no committed reward floor; post-pool-exhaustion income undefined (C-05).
- **Solvency condition:** measurable-unit launch gate (EC-3), a bootstrap compensation mechanism that a professional operator can underwrite (EC-3), early-epoch emission weighting made concrete, and a published fee-market successor model for the post-emission era (EC-1).

7.3 Issuer (KYC organisation)

- **Inflows:** mint fees, leases, disclosure brokerage, its own client-side KYC pricing — all undefined.
- **Outflows / risks:** ≈2,880 root transactions/day forever (§6.3); AML per-query fees; vault and broker operations; slashable bond of unpriced deterrent value; auto-suspension on bond exhaustion (deterministic, no grace); legal exposure as sole identity-link holder; H5/L2 counsel outcome risk.
- **Solvency condition:** an issuer business case with break-even portfolio size at ratified fee bands (EC-1); relayer/root-cost incidence defined (who pays the 30-second cadence, at what gas price); bond sized in measurable units (EC-3).

7.4 Principal (human/organisation deploying agents)

- **Inflows:** disclosure bids (optional, price-floor-protected); the amortisation saving (one KYC, N agents).
- **Outflows / risks:** KYC cost; mint + lease per agent; capital F × 3–12 months at zero yield per agent (no pooling in committed scope); freeze risk through disputes incl. judicial phase; ladder loser-pays exposure; NC1 blacklist and NC4 depeg tail risks on locked collateral.
- **Solvency condition:** published F bands and TTLs calibrated so that locked capital is proportionate to segment value (H4 calibration); a griefing-cost analysis showing dispute-freeze attacks are unprofitable (EC-6-adjacent, see M-06 remediation).

7.5 Relying party / verifier (incl. x402 sellers and facilitators)

- **Inflows:** free verification (by construction); funded recourse up to F; optional paid disclosure.
- **Outflows / risks:** claim bonds (JIT or permanent); recourse capped at F with band values unpublished (cannot price residual risk); S-7 unresolved (a seller that never registered has no standing — or a facilitator fronts bonds and becomes an unregulated de-facto insurer, a council matter the corpus flags).
- **Solvency condition:** published F bands; S-7/D-x3 decided; bonded-tier value proposition repriced or merged (M-07).

7.6 Arbiter

- **Inflows:** loser-borne fees only. **Risks:** bonded capital idle at zero dispute volume; KYC/registration overhead.
- **Solvency condition:** an explicit bootstrap mechanism — treasury-funded retainer from the remediation/treasury compartment, minimum-fee guarantees, or delayed activation of the F2+ rail until organic volume exists (EC-6).

7.7 Treasury / foundation (steward)

- **Inflows:** genesis operations allocation (size ambiguous — C-01); slash remediation shares (burn-only if unclaimed — correctly conflict-proofed); optional protocol rent (undefined).
- **Outflows / risks:** the entire unbudgeted stack (H-03): two NC2 contract audits, the Cluster-A ZK audit package (four circuits audited together; market reference >$150k for ZK scopes), the R20 slate (two chain audits + crypto audit + pentest + economic review), the independent simulation, multi-jurisdiction counsel (≥6 workstreams), trademark clearance, foundation + SL incorporation and operation, an independent co-signer custodian in perpetuity, and a **TVL-scaled bug bounty that is a standing liability growing with adoption** and has no funding source.
- **Solvency condition:** a published fiat budget and runway statement naming the funding source (EC-5). The implicit answer — the CONECTIA OÜ balance sheet — is never written down, and is itself undisclosed.

8. Survivability and runway

8.1 The dependency graph to mainnet (as documented)

[D] Hard gates before mainnet-with-funds: C2-P0 circuit audit; NC2 escrow-contract audit(s); NC4 depeg parameters; H5/D.1 council mechanics (counsel); H2 dispute protocol + arbiter market; consent governance (counsel); M6 cross-border liability; R20 audit slate; genesis prerequisites (value freeze + simulation) (Agent-KYC-Roadmap.md gate table). The roadmap itself declares the two longest-lead items — the independent simulation and the coordinated counsel engagement — "critical-path… every month they remain uncommissioned is a month added to mainnet", and neither was commissioned at freeze.

[A] Circularities embedded in that graph (Finding H-05): 1. Fees/bands/tokenomics ⇄ simulation ⇄ pilot calibration (H4/M3/M5) ⇄ live product ⇄ audits ⇄ genesis ⇄ simulation. 2. Sybil floor ⇄ x402 adoption economics (S-8/D-x5) — both assigned to the same uncommissioned study. 3. F2+ dispute rail ⇄ arbiter market existence (H-02). 4. Bonded-verifier value ⇄ dispute-rail utility ⇄ bonded-verifier population. 5. Cluster-A audit couples four circuits — one slip blocks the whole F2+ economy.

These loops are breakable — by commissioning the simulation with explicitly scenario-based (not pilot-dependent) v1 parameters, by staging launch without F2+ (identity-only MVP at MA3 needs none of the blocked gates, as the x402 study correctly certifies), and by an arbiter bootstrap subsidy — but each break must be decided, not discovered.

8.2 Implementation plan realism

[D] The Phase-0 backlog starts from a scaffolding-only repository, defines ~90 cards across 8 epics with S/M/L sizing only — no hours, FTEs, dates or costs — and is **desynchronised from the frozen specs**: card E0-3 encodes a nullifier design SPEC-2 explicitly rejects as "a category error"; no cards exist for k_release/R-A13.1 (launch-blocking), the §12.5 session-reservation model, the R-A28 arbiter registry, or the D-x5 fee study; G0-3 books the "independent economic simulation" as an internal Python tools card while the roadmap specifies an external engagement with the project's longest lead time (M-04).

[J] The plan demonstrates scope completeness discipline but contains no schedule, no budget and at least four launch-blocking work items with zero cards. As an implementation plan for investors it is a work-breakdown structure, not a plan.

8.3 Tail risks on the collateral (NC register)

[D] NC1 (USDC/EURC issuer blacklist): "accepted + documented" — an uninsured, unpriced loss scenario across all locked escrow. NC4 (depeg): parameters (1.05× overcollateralisation, >5%-for->24h degradation, per-asset haircuts) drafted but **not ratified** (gate MINT-138). NC2 (contract exploit): audit-gated, bounty-backed. NC5 (AML provenance): per-mint operating cost. NC6 (multi-collateral governance): standing governance overhead; BTC rail adds volatile-collateral haircut risk; XMR rail is counsel-gated with a real possibility of never activating (AMLR 2024/1624).

[J] Individually reasonable; collectively they mean the "funded recourse" guarantee carries issuer-censorship and depeg tails that the relying party cannot see in the escrow_coverage_gte predicate. Disclosure language for verifiers should state that F is denominated in stablecoin claims, not risk-free value.

8.4 Competitive window

[D/External] The corpus's timing thesis — that transparent bearer identity could standardise before MintID ships — is confirmed and accelerating: KYA primitives at every major payment network (April 2026), Skyfire × Visa Intelligent Commerce, F5 enterprise distribution, ERC-8004 anchoring, Concordium live with merchants. MintID's counter (the consented commercial rail as KYA-parity wedge, testnet-first demo) is correctly sequenced in the roadmap but sits behind the uncommissioned critical-path engagements.

[J] The realistic strategic read: MintID will not win the low-assurance, card-adjacent mass segment — that is already taken. Its defensible niche (privacy-preserving, funded, regulated-grade backing) remains unoccupied and externally validated as unserved. The survivability question is therefore almost entirely **execution-side** (funding, sequencing, the five critical findings), not demand-side.

8.5 Scenario view

Scenario	Assumptions	Outcome for stakeholders
Base	C-01…C-04 remediated in ≤ 1 quarter; simulation + counsel commissioned now; identity-only MVP (MA3) ships against testnet while F2+ gates close; genesis ≥ 12 months out	Viable path; validator and issuer solvency determined by simulation outputs; token case rests on ratified cut/burn bands
Adverse	Critical findings unremediated; genesis attempted on current documents	Gate unmeasurable (C-02) → either gate waived (security theatre) or launch stalls; validator recruitment fails; development entity potentially under-capitalised 11–23× (C-01); audit stack unfunded (H-03)
Competitive-loss	Execution correct but slow; transparent rails capture the relational segment and add "good-enough" privacy features	MintID relegated to a compliance niche; five-flow revenue shrinks to flow-B scale; token security budget insufficient post-emission (C-05 consequence)

---

9. Findings register

Severity scale: **Critical** — blocks a favourable viability opinion; must close before genesis parameters freeze. **High** — material threat to solvency or survivability; must close before mainnet. **Medium** — should close before mainnet; acceptable to schedule. Each finding states evidence, impact and the **required remediation (claim/correction)**.

Critical

C-01 · Genesis allocation band contradiction (irreversible) - *Evidence:* Genesis-Allocation-and-Transparency-Policy.md §2 L28, §2.1 L30 — 10–20% of MAX_SUPPLY band (10.8M–21.6M tokens) vs working figure ~1,000,000 tokens with the in-text illustration "1M ≈ 0.93% — conservative within the band"; §2.2 L31 — no post-genesis mint. - *Impact:* an 11–23× ambiguity on the development entity's permanent capitalisation; both external communications and the simulation brief inherit the ambiguity. - *Required remediation:* issue a corrected policy revision that either (a) restates the band as "≤ 20%, with a working point near 1%" and justifies the low point against the H-03 cost stack, or (b) adopts a figure inside the stated band with dilution disclosure. Record the decision with rationale, as was done for MINT-150.

C-02 · Circular launch dependency (no price → unmeasurable gate → unrecruitable validators) - *Evidence:* SPEC-1 §5.4.1 L303 (gate denominated in bonded-stake *value*); Genesis Policy §3 L45 (no listing), §6.5 L94 (MiCA path post-genesis), §1.6 L22 (retroactive-only grants); SPEC-1 §6.2 L366 (min_bond 1000 tokens placeholder). - *Impact:* the launch gate cannot be evaluated; the issuer bond has no known deterrent value; professional validators have no underwritable compensation; the entire §5.4 apparatus — the project's own answer to the young-chain attack problem — is non-operational as specified. - *Required remediation:* (i) redenominate all launch thresholds in measurable units (bonded % of supply, independent-operator count, Nakamoto coefficient, fiat-denominated attack-cost proxy with published methodology); (ii) define a compliant validator-bootstrap compensation instrument (e.g. service agreements in fiat/stablecoin from the operating budget, which the genesis-

allocation exemption analysis already permits for contractors) instead of relying on unpromisable grants; (iii) size min_bond in the same measurable units.

C-03 · Ghost normative requirements R-A13/R-A14/R-A15 - *Evidence:* SPEC-2 cites the series ≥10× (:5, :19, :63, :65, :358, :432) and defines only R-A13.1; the only extant text is archive/Agent-KYC-Concept-and-Market-Paper-v2.md :410–412, where R-A14 mandates a bonded fiat custodian — the model the Jun 2026 pivot abolished; "R-A15" is used with three incompatible meanings (remediation ladder / freeze precedence / cover-and-keep-private council mechanics — gate H5/MINT-82). - *Impact:* the funded-accountability layer — the product's core promise — has no authoritative normative definition; auditors, issuers and counsel following the trace land on superseded text; gate H5's subject ("R-A15 normative text") is ambiguous. - *Required remediation:* write R-A13/R-A14/R-A15 into SPEC-2 as current normative text (self-custodied escrow, three-exit rule, cascade), renumber or disambiguate the three "R-A15" usages, and mark the archive text superseded in place.

C-04 · No ratified economic parameters; critical-path simulation uncommissioned - *Evidence:* SPEC-1 L324, L1250–1263 (all economic parameters open); Genesis Policy §6.2 L91 ("Phase 0: commission the economic simulation"); Agent-KYC-Roadmap.md :38–41 ("every month they remain uncommissioned is a month added to mainnet"); backlog G0-3 books it as an internal tools card (M-04). - *Impact:* nothing downstream — genesis, fees, bands, bonds, emission, gate thresholds — can move; the project's clock is running against an engagement that does not exist. - *Required remediation:* commission the simulation now as an external engagement with an investor-grade statement of work that includes: a **protocol P&L under published adoption scenarios**, flow-A/C protocol-cut and burn-split bands, the §5.4 deliverables, the D-x5/S-8 x402 fleet scenarios, and explicit v1 parameters that do not depend on pilot data (breaking loop H-05.1).

C-05 · Structurally bounded protocol revenue and token value accrual; undefined post-emission security budget - *Evidence:* SPEC-1 R15/R18/§10.5 L806 (verification free, presentations off-chain), §9.3 L650–662 (state O(registries), not O(population)), R26 L77 + §14.4 L1076 (premium services barred from consensus rewards, no protocol cut), §5.2 L266–283 (validator income = emission only), L472 + L1263 (agent mint/lease non-normative, parameterless); SPEC-3 §8/R-A28 (arbiter fees off-chain, same bar); whitepaper §8 (protocol cuts "partially burned", percentages undefined). - *Impact:* token demand and validator budget do not scale with adoption; at reward-pool exhaustion the security budget's successor is undocumented; the token investment case reduces to blank percentages. - *Required remediation:* (i) make agent mint/lease economics normative with ratified band ranges; (ii) ratify non-zero protocol-cut/burn bands for flows A and C; (iii) require the simulation to model the post-emission fee market and state the intended successor (fee market sizing, protocol rent activation criteria, or an explicit decision to accept a shrinking budget with rented-security fallback).

High

H-01 · Adverse unit economics in the flagship x402 segment - *Evidence:* Use-Case-x402 §2.3 :47 (externally corroborated; avg tx $0.20–$0.30); SPEC-2 :360, :367 (per-agent prefunding, CCTP cycle); S-8/D-x5 unresolved. - *Remediation:* publish the unit-economics study (mint + lease + gas + locked-capital opportunity cost vs per-agent value by segment) **before** the MA3 go/no-go; define the minimum viable segment explicitly.

H-02 · Arbiter market cold-start - *Evidence:* SPEC-3 R-A28 :81 (launch-blocking, loser-funded only). - *Remediation:* fund an arbiter bootstrap (retainers from operations treasury, minimum-fee guarantees, or staged F2+ activation on organic dispute volume); add the missing R-A28 implementation card (M-04).

H-03 · Unbudgeted cost stack, no documented fiat funding source - *Evidence:* R20 slate (SPEC-1 Phase 6 L1192–1198); NC2 dual audits, Cluster-A ZK audit, TVL-scaled bounty (SPEC-2 :366, :406, :443); counsel engagement list (Agent-KYC-Roadmap.md :41); foundation + SL + independent co-signer (Genesis Policy §4–§5); external reference costs: ZK audit scopes routinely >$150k, protocol audits $25k–$100k each, remediation passes $5k–$20k. - *Remediation:* publish a fiat budget and runway statement naming the funding source and its sufficiency through mainnet + 12 months; include the bounty as a TVL-indexed contingent liability.

H-04 · Genesis accounts unrestricted in delegation/voting - *Evidence:* Genesis Policy §4.2 L58, §4.4 L66 (transfer restrictions only); Cosmos vesting accounts can delegate; SPEC-1 §5.4.1 thresholds. - *Remediation:* adopt an explicit no-delegation/no-vote rule for genesis allocation accounts, or exclude them (and delegations traceable to them) from launch-gate arithmetic; publish the rule in the Genesis Allocation Statement.

H-05 · Circular calibration dependencies - *Evidence:* §8.1 of this report; Agent-KYC-Roadmap.md :73 (H4/M3/M5 "pilot-dependent"); gate table. - *Remediation:* the simulation SoW must deliver scenario-based v1 parameters explicitly not gated on pilot data; adopt the staged launch (identity-only MVP first) as the official sequencing, which the x402 study already supports.

H-06 · Issuer concentration and related-party economics - *Evidence:* §6.3 of this report; SPEC-2 R9/R-A6.1/:443; Genesis Policy §5 L78–83; whitepaper §14.3. - *Remediation:* model issuer break-even in the simulation; define relayer/root-cost incidence; add a related-party disclosure section (first issuer, R26 operator, development company, steward) to the investor pack and transparency report.

H-07 · Competitive timing (externally confirmed) - *Evidence:* §4 rows 4, 8; whitepaper §2.4/§13.3. - *Remediation:* protect the already-planned testnet-first commercial-rail demo from the critical-path slip by decoupling it from gated F2+ work (it needs none — x402 study §4.4); treat the uncommissioned engagements (C-04) as the timing risk's binding constraint.

Medium

M-01 · Rented security unbudgeted and in scope tension — ICS requires IBC at genesis vs "interchain connectivity after mainnet stability" (Could-have) and the bridges Won't-have; ICS/Babylon payment flows appear nowhere in emission or reward accounting. *Remediation:* the §5.4.3 evaluation must include cost, scope-conflict resolution and an emission-accounting amendment if adopted. **M-02 · Reward-pool and fee-module inconsistencies** — "other uses of the reward pool" (§5.4.4 L306) undefined and in tension with Genesis §2.2; x/fee described as conditional (L470) but normatively required (L371). *Remediation:* editorial reconciliation pass on SPEC-1 §5–6. **M-03 · R20 sequencing gap** — economic values freeze at Phase 0 (L1123) but the independent economic review sits at Phase 6 (L1196) with no unfreeze mechanism if it objects; Genesis §4.5/§6.3 imply pre-genesis audit. *Remediation:* state the unfreeze/re-freeze procedure and align the two documents' sequencing. **M-04 · Phase-0 backlog desynchronised from frozen specs** — E0-3 encodes the rejected nullifier design ("category error", SPEC-2 :371); no cards for R-A13.1/k_release, §12.5 reservations, R-A28 registry, D-x5 study; G0-3 mislabelled internal. *Remediation:* backlog reconciliation sprint before any implementation contract is let against it. **M-05 · Unpriced collateral tail risks** — NC1 "accepted + documented" (uninsured blacklist loss on TVL); NC4 parameters unratified (MINT-138). *Remediation:* ratify NC4 parameters; add verifier-facing disclosure that F is a stablecoin claim subject to NC1/NC4 tails. **M-06 · Economic griefing vectors** — capital-freeze attacks via contested claims (freeze survives expiry/self-revocation through the judicial phase); "cheap to be rich" ladder dynamics; bond sizing open (SPEC-3 :399). *Remediation:* bond-sizing study with griefing-cost floor (attacker's expected loss > victim's freeze cost), freeze-duration caps per rung (partially present — :200), and a counter-bond assistance mechanism for capital-constrained honest principals. **M-07 · Bonded-verifier tier economically dominated** — same verification, same rail access via JIT bond, lower carrying cost permissionless. *Remediation:* price the bonded tier's privileges explicitly (latency, limits, standing) or collapse to one tier + JIT. **M-08 · Renewal ⇄ re-disclosure conflict of interest** — the issuer earns brokerage on every re-identification that credential churn (3–12 months, R-A29) forces; the same side influences grade validity cadence. *Remediation:* disclose; consider capping brokerage on renewal-driven re-reveals or fixing cadence per grade in x/identityparams (it largely is — keep issuer discretion out of it). **M-09 · Uncorroborated insurance-market figure** — $4.48B→$25B (SPEC-3 :414) vs independent $0.3B–$1.2B (2025). Deferred-extension scope only. *Remediation:* re-source before any partner/investor use; the deferred insurance extension's demand case should not inherit it. **M-10 · Documentation hygiene** — Genesis header v1.1 vs footer v1.0; utestbit pre-rebrand denomination; public website acknowledged stale (MINT-127); cap change 93,924,792→108,000,000 without recorded economic rationale; repository publishes working figures the website policy says are unpublished (channel inconsistency relevant to MiCA marketing-communication discipline). *Remediation:* single editorial pass + a rationale memo for MINT-150.

---

10. Conditions precedent to a favourable opinion

The auditor would upgrade to an unqualified viability opinion when the following are evidenced. EC-1…EC-10 are the audit's claims/corrections; the right column maps to the project's own gate table where one exists.

#	Condition precedent	Closes findings	Project gate mapping
EC-1	Independent economic simulation commissioned (external SoW) delivering: protocol P&L under published adoption scenarios; flow-A/C cut and burn bands (non-zero, ratified); agent mint/lease normative bands; issuer break-even model; post-emission security-budget successor; D-x5/S-8 x402 fleet scenarios; scenario-based v1 parameters not gated on pilot	C-04, C-05, H-01, H-05, H-06	MINT-79/80 (G0-2/G0-3), D-x5
EC-2	Genesis allocation contradiction resolved by recorded, justified decision	C-01	Genesis Policy rev; R20 scope
EC-3	Launch gate, issuer min_bond and validator bootstrap redenominated in measurable units; compliant validator compensation instrument defined	C-02	SPEC-1 §5.4; Genesis §1.6
EC-4	No-delegation/no-vote rule (or gate-arithmetic exclusion) for genesis accounts, published in the Genesis Allocation Statement	H-04	Genesis §4.1
EC-5	Fiat budget and runway statement: full cost stack (audits, counsel, simulation, foundation, custodian, bounty-as-contingent-liability) with named funding source through mainnet + 12 months	H-03	— (new)
EC-6	Arbiter-market bootstrap mechanism funded and documented; R-A28 implementation card exists	H-02	H2 gate (MINT-84/94/140)
EC-7	x402 unit-economics study published before the MA3 go/no-go	H-01	MA3 exit criteria
EC-8	R-A13/R-A14/R-A15 written as live normative text; “R-A15” disambiguated	C-03	H5 gate (MINT-82) precondition
EC-9	NC4 parameters ratified; NC1/NC4 verifier-facing risk disclosure adopted	M-05	MINT-138

#	Condition precedent	Closes findings	Project gate mapping
EC-10	Related-party economics disclosure (first issuer / R26 operator / development company / steward) in investor pack and transparency report	H-06	Genesis §4.6 extension
EC-11	Backlog reconciliation with frozen specs (E0-3 corrected; missing launch-blocking cards added)	M-04	Phase-0 plan rev
EC-12	Editorial reconciliation pass (M-02, M-03, M-10 items)	M-02/03/10	—

Of these, **EC-2, EC-3, EC-4, EC-8, EC-10 and EC-12 are documentation/design decisions executable in weeks at negligible cost**. EC-1 and EC-5 are the two spends that unblock everything else — and the corpus itself already says so.

---

11. Acknowledged strengths (practices to preserve)

1. **"Mechanism committed, numbers gated"** — the refusal to publish invented figures, with house rules requiring primary sources for every market number (whitepaper §8.1). This audit's own validation confirms the cited figures were chosen honestly.
2. **Hard-cap discipline with no re-mint from burns**, no perpetual-deflation promises, and no stake-tax gimmicks (SPEC-1 §6.1).
3. **Conflict-proofed remediation**: unclaimed slash remediation can only be claimed late or burned — "the body that adjudicates slashes must never profit from unclaimed remediation" (§6.2 L371).
4. **Native, consensus-enforced vesting** (4y/1y cliff) + labelled genesis addresses + hash-committed allocation statement + independent co-signer custody (Genesis §4).
5. **Prudent MiCA posture**: no listing before the deliberate Title-II path; no voluntary MiCA-style white paper; worthless-by-construction testnet tokens (Genesis §1, §3).
6. **First-implementer without privileges**, verifiable on-chain (same bond, same admission, same roots — whitepaper §14.3).
7. **The §5.4 launch gate held to the same rank as the audit slate** — the right instinct; it needs only the C-02 redenomination to become operational.
8. **The escrow pivot** to self-custodied USDC/EURC with the three-exit rule and k_release separation — eliminating fiat custody, chargebacks and the custodian honeypot, and closing the leaked-secret blast radius on funds (SPEC-2 §12.6).
9. **The right pilot metric already identified**: share of disputes closed at L0 with zero identity revealed (whitepaper §8.1) — the direct measure of the money-before-identity thesis.
10. **Honest self-registers**: the weak-points lineage, the gate table ("the card is the gate"), dated claims flagged for refresh, and feasibility verdicts that include "Not acceptable" rows against the project's own ideas.

---

12. Appendices

Appendix A — Complete inventory of concrete figures in the corpus

Figure	Value	Source
MAX_SUPPLY	108,000,000 (immutable)	SPEC-1 §6.1 L315; MINT-150
Superseded cap	93,924,792 (2026-07-12)	SPEC-1 L8; MINT-118
Dev allocation working figure	~1,000,000 tokens (≈0.93%)	Genesis §2 L28, §2.1 L30
Declared team/dev band	10–20% of MAX_SUPPLY	Genesis §2.1 L30
Vesting	4 years continuous, 1-year cliff	Genesis §4.2 L58
Post-genesis lock-up	12 months default	Genesis §4.4 L66
Issuer min_bond (placeholder)	≈1000 tokens, provisional exponent	SPEC-1 §6.2 L366
Remediation TTL	~1 year (governed)	SPEC-1 §6.2 L371
Active validator set	100 (candidate)	SPEC-1 §5.1 L256
Block time / finality	2–3 s / p95 ≤ 5 s	SPEC-1 §5.1 L257–258
Status heartbeat / max root age	30 s / 45 s fail-closed	SPEC-1 §9.1 L594–596
Presentation lifetime	10 s (R14)	SPEC-1 L62
Grade validity	A1 12m · A2 12m · A3 6m · A4 3m	SPEC-1 §4 L228–233
Council threshold	5-of-9 + independent custodian	SPEC-1 §7.1 L399
Rented-security window	first 12–24 months	SPEC-1 §5.4.3 L305
Session budgets	Moment A ~6 s (cap 12 s); B ≤ 10 s; illustrative ~120/300 s	SPEC-2 :397–404
Reservation TTL default	~1 h (illustrative)	SPEC-2 :402
CCTP funding cycle	~19–21 min per funding event	SPEC-2 :367
Credential renewal cadence	3–12 months	whitepaper R-A29
Depeg parameters (unratified)	1.05× overcollateralisation; >5% for >24 h	Roadmap NC4/MINT-138
x402 (late Apr 2026, dated)	~69k agents; ~165M tx; ~$50M cumulative	Use-Case-x402 §2.3 :47
MiCA exemption thresholds	<150 persons/Member State; ≤€1M/12mo	Genesis §3 L41
Corporate structure	Spanish SL, 100% CONECTIA OÜ; MintID foundation as steward	Genesis §5 L78–83
Phase-0 plan volume	8 epics, ~90 cards, S/M/L only	Backlog :187

(This table is exhaustive for economic figures: no fee, band value, protocol cut, burn split, emission parameter or bond size exists anywhere in the corpus beyond the rows above.)

Appendix B — Auditor derivations used in this report

Derivation	Inputs	Result
x402 avg value/tx	$50M / 165M tx	≈ $0.30 (externally confirmed $0.20–$0.30)
x402 tx/agent	165M / 69k	≈ 2,391
x402 lifetime volume/agent	$50M / 69k	≈ $725
Issuer root cadence	86,400 s ÷ 30 s	2,880 tx/day/issuer
Genesis band contradiction	10–20% of 108M vs 1M	10.8M–21.6M vs 1M → 11–23×
min_bond share of supply	1000 / 108M	0.000926%

Appendix C — External sources (accessed July 2026)

Listed inline in §4. Primary corpus sources cited throughout as file:section/line.

Appendix D — Auditor's note on independence and related-party disclosure

This audit was prepared by **Conectia PRO** from the frozen document corpus and public sources only. The auditor has no position in the project's token (none exists).

Related-party disclosure. Conectia PRO is not an arm's-length third party: its founder is also the principal of the MintID project, and the corpus identifies the CONECTIA group (CONECTIA OÜ / Spanish SL) as the development company and prospective first implementer (Finding H-06). Conectia PRO has no commercial interests beyond providing services to MintID. Independence in this engagement is therefore **methodological rather than corporate**: every material claim is cited to the frozen corpus or to public sources; auditor arithmetic and judgment are labelled and separable (**[D]/[A]/[J]**); and each finding is drafted so that a third party can re-perform it from the citations alone. This disclosure applies to the auditor the same standard the report imposes on the project (EC-10), and it does not substitute for the independent economic review the corpus itself mandates (R20; EC-1/EC-2): where due diligence requires an arm's-length opinion, this report should be read as a rigorous, fully-cited self-assessment.

The findings register is intended to be attachable as-is to investor or partner due-diligence packs; the conditions-precedent matrix (§10) is drafted so that each condition is objectively evidenceable.

---

*End of report. Prepared by **Conectia PRO** during July 2026; delivered 2026-07-30. Qualified opinion — see §2.1. Contact the engagement owner for the underlying citation extracts.*