

What the AO Counters Do Not Measure

A data-feasibility study on Arweave and AO, and why the public activity metrics cannot yet support a valuation

Nam | 15 September 2026

Disclosure: I hold personal positions in both AR and AO. I defined the measurement framework and the decision rules before analysing the data, and the study permitted four outcomes: both, AR only, AO only, or neither. It reached none of them. Every empirical figure below comes from public endpoints I queried myself; raw responses, queries, timestamps and checksums are preserved.

I set out to underwrite a paired position in AR and AO. The thesis was that AO turns Arweave's permanent message history into live computation, and that growing AO activity would therefore show up as measurable storage demand for Arweave.

I could not substantiate it. Not because the thesis is wrong, but because the metrics everyone cites — messages, processes, protocol fees — do not measure what they are assumed to measure. This is a report on why, and on what would have to be measured instead.

Three claims I made early in the work did not survive contact with the data. I retract them below.

/ The process spike was one issuer, not a thousand applications

On 25 July the network counters recorded 2,305 new processes against 679 messages in the same rolling window. The obvious reading is mass deployment followed by abandonment. I queried the block range directly to check.

Two case-sensitive tag schemas over blocks 1,965,241–1,965,907 returned 7,054 unique IDs with no overlap between them and no unresolved pagination. Of those, 2,152 were process creations.




Figure 1. The spike has one author. 2,078 of 2,152 processes ran name-token@1.0; 2,084 were created by a single owner address. Source: arweave.net GraphQL, queries and responses preserved.

These are permaweb name tokens, not applications. The Initial-Owner field spans 416 distinct addresses, but the creating owner is one. Read together with Arweave's published consolidation of legacy naming systems, a bulk issuance or migration event is the more plausible explanation than a thousand product launches.

The consequence is directional. A name token can be resolved thousands of times without ever producing a state-changing message. **Low message counts do not indicate a dead process here; they indicate that message count is the wrong instrument for this object class.** Process creation is not an application count, and the ratio of messages to new processes mixes populations that should not be divided by one another.

/ Message activity is similarly concentrated

The same window contained 1,147 messages, addressed to 43 distinct recipients.




Figure 2. One process received 58.2% of messages. The top three recipients account for 70.5%. Dominant actions in the window — PlaceDecorations, Move, WaterMultiple, PlantSeeds — are consistent with a single game-like application rather than diversified demand. Source: same queries as Figure 1; 154 messages had no recipient tag.

This does not mean the activity is fake. It means that in this window, aggregate message counts describe the behaviour of one or two applications, and cannot be read as a measure of ecosystem breadth.

/ The fee lives in the bundle, not in the message

The economic link from AO to Arweave is the one that matters for a paired position, and it is the hardest to trace. AO DataItems are written to Arweave inside ANS-104 bundles. The individual item carries no fee; the parent bundle transaction pays the L1 cost.

I downloaded one parent bundle, parsed its header, and reconciled every child against the index.




Figure 3. Anatomy of one bundle. 7,092 = 352 (shared header) + 1,250 + 1,156 + 1,545 + 1,244 + 1,545. Five DataItems: three Messages, two Assignments; four ao.N.1 and one ao.TN.1. All five child IDs were found in the index with matching parent links. Source: arweave.net/raw and GraphQL; binary and audit preserved.

Two observations follow, and they point in opposite directions.

First, fee attribution is technically possible. Allocating the parent fee by serialized bytes is a defensible convention. But it is a convention, not a protocol field: it distributes the shared 352-byte header proportionally, it is not a marginal cost, and it cannot be multiplied across the network from one conveniently chosen bundle.

Second, and more important: the combined body data across all five items is **six bytes**. Almost the entire 7,092 bytes is signatures, tags and framing. Any model that estimates AO's storage footprint from payload size will understate it by orders of magnitude.

The L1 fee payer is a known address. Whether the end customer reimbursed it is not known. ANS-104 explicitly permits a third party to pay for a DataItem it did not sign, and AO's own bundler documentation describes a subsidy path. **Both of these can be true at once: real AR cost was incurred on L1, and no end customer paid for it.** A subsidy does not erase the transaction, but it changes the forecast for how durable that demand is.

/ The reconciliation that failed

I could not make the aggregate counters agree with the record-level queries, and I am reporting the gap rather than resolving it by assumption.

Same block range	Record-level query	Counter increment	Gap
Messages		1,147	677 +470
Processes		2,152	2,305 −153

Both queries terminated with `hasNextPage = false` and no duplicate IDs. Completed pagination is not the same as reconciled data. Candidate explanations — counter definitions, index coverage, retroactive updates, tag mapping, or the documented possibility that `DataItems` change block-height association on bundler retry — were not narrowed to one.

This is a control that did not pass. Until the denominator is understood, no statistic resting on it should be published as a network-wide rate.

/ What I retract, and what remains unmeasured

Three earlier readings are withdrawn. **The 1.6 billion legacy versus 171,743 mainnet message split cannot be used as an adoption argument** — the official AO token integration guide still uses `MODE: 'legacy'`, and a recent sample contains an `ao.TN.1` transfer addressed to the AO token process itself. Legacy is a message format, not a verdict on economic value. **The process spike is not evidence of abandonment. And a bundle's L1 fee is not proof of customer revenue.**

Metric	Status after this study
Process creation, classified by device and owner	Measured, within a defined index scope
Serialized bytes and parent-bundle fee	Measured on one sample
Active processes, 30-90 day retention	Not publishable — cohorts under-aged, execution paths incomplete
Fee-paying end customers	Not measurable from public data
AO service revenue, sustainable payment flow	Not demonstrated
Valuation of AR or AO	No basis from this work

I sampled twenty processes to test whether retention could be tracked, selected by hashing IDs before querying so the sample could not be chosen on outcome. The queries completed. I did not publish a retention rate, because a rate computed over an unreconciled denominator with incomplete execution coverage would be a number with no meaning.

/ Why I still hold both

None of this makes AR or AO uninvestable. It makes them unmeasurable with public data today, which is a different statement.

My position rests on an architectural bet, not on the metrics above, and I would rather say so plainly than dress it in numbers that do not support it. The bet is that software with long-lived state needs a canonical history that outlives the operator who created it — that a process should be able to survive its front end, its company and its compute provider. Arweave stores that history; AO interprets it into live state. If that pairing matters, it will eventually show up as payments. It has not yet.

What would change my mind, in either direction, is specific:

- **A reconciled denominator.** Counter definitions and record-level queries agreeing over a bounded range. Without it, every rate built on top is unverifiable.
- **One traceable payment chain.** Customer → successful settlement → provider → L1 cost or tokens encumbered. One chain, fully traced, is worth more than another million rows from the same filter.
- **Retention measured on aged cohorts**, defined as a successful fee-bearing execution in a specified age window, with automation and configuration traffic separated out.
- **Evidence on who funds the bundler.** If AO's storage footprint is subsidised, the demand is real but its durability is not established.

If those four resolve favourably, the paired thesis becomes underwritable. If the payment chain cannot be traced at all, the honest conclusion is that AO may succeed as software without materially repricing AR — and I would rather find that out from data than from price.

/ Method, scope and limits

Scope. Cut-off 15 September 2026 UTC. Aggregate snapshot generated 14 September 00:00:21 UTC at height 2,000,384, with 59 historical observations from 17 July (30 August missing — the series is not 59 contiguous UTC days). Record-level window: blocks 1,965,241–1,965,907 inclusive, two tag schemas, 7,054 unique IDs, pagination complete. Bundle sample: parent `tTDAMI7Gj fmpQs8Ma- aEZBwyMxD4mpRBdIsare0Js0U` at height 2,001,020. No wallet was connected, no transaction signed, no payment made, and the project team was not contacted.

Known limits. Rolling counter fields use a 720-block window, not a UTC day, and must not be summed as disjoint days; the final snapshot's window covers 95,152 seconds. Counter changes are provider-reported, not a reconciled chain census. Query completeness means completeness within the index at retrieval time, not that all protocol activity is indexed. Fee allocation by serialized bytes is an illustrative convention. `DataItem` signatures, block proofs and snapshot commitments were not independently verified. An earlier interrupted bulk download (87,646 rows, 87,599 unique IDs, 58 of 164 planned files) is retained as audit evidence and was not used for any statistic here.

Sources. Arweave GraphQL (`arweave.net/graphql`) and `arweave.net/raw`; AO network metrics process snapshot; ANS-104 bundle specification (ArweaveTeam/arweave-standards); AO token mainnet integration, MU/bundler and GraphQL-indexer FAQs (`permaweb/ao-token-docs`); AO posts on Permaweb Names, Open Service Markets, NASA, Decentralizing the Financial Layer and Bazar. Queries, raw responses, retrieval timestamps, derived calculations and SHA-256 checksums are preserved in the evidence archive. Checksums detect post-packaging modification; they do not prove the indexer or the chain.