

IAM — Intelligent Asset Management

The analytical layer of MetaNetwork EQI. It supplies the analysis on which the network is governed, and deterministically executes the protocol rules that keep the active layer invariant.

An analytical model, not an agent: it narrows scenarios and ranks them by probability – judgement stays with the DAO.

v1.0

Intelligent Asset Management

Patent pending (PPA filed)

0 · MANDATE

Two circles of authority

IAM is the analytical organ of the meta-network. Its authority is split deliberately, and the split is architectural — not a policy that can be reconfigured.

Analysis — no decisions

IAM collects, verifies and analyses data, then narrows the field to 1–3 justified variants ranked by probability, each with sources, grounds, limits and error margin. It issues **no recommendations**. Judgement belongs to the DAO.

Execution — deterministic rules only

Where the outcome is already prescribed by the protocol, IAM executes it: α - redistribution of the active layer, and **STOP** on any deviation of A_active from the constant. It applies a rule — it does not choose.

Where there is judgement – the DAO decides. Where there is a formula – IAM executes. IAM cannot exceed its competence, because it has no channel to do so.

A different class of system

Not an LLM – the core is not a transformer


Not an agent – no autonomous action channel

Not an oracle – internal analysis, not external attestation

Not a judge – it does not rule on ownership or fairness

An LLM may be used inside IAM as one tool among many — for instance to phrase a conclusion in natural language for the DAO. That is peripheral. The analytical core is not an LLM, and LLM output is treated as just another source that must pass verification.

Narrowing, not expansion



A transformer expands the space of continuations. IAM narrows the field of variants to a consistent calculated probability.

A transformer takes a query and opens context around it — words, links, associations — then selects the most likely continuation from that expanded space. IAM does the opposite: it takes all available data, forms the field of possible system states, discards the weakly-supported ones, and outputs the few that survive. Not expansion — concentration.

Why a transformer cannot do this work

- **Self-attention zeroes out weight.** Irrelevant tokens are effectively filtered away. For analysing a live system that is a defect: every event, every anomaly, every state carries weight and may not be nulled. IAM keeps all verified data — the difference between items is depth of analysis, not exclusion.
- **Uniform stacked blocks.** Every transformer layer performs the same operation. IAM instead runs a sequence of functionally distinct steps: ingest, verify, sort by depth, multi-stream processing, output construction.
- **Softmax over a vocabulary.** A transformer's output is a probability distribution over tokens; any "analysis" is behaviour simulated on top of token prediction. A structured analytical output — sources, probabilities, grounds — does not exist there as a primitive. In IAM it is the primitive.

IAM does not work with text tokens. It works with structured real-world signals: timestamped event series, network and account states, causal dependency graphs, market data, anomaly markers, $\alpha/\beta/\gamma$ stabilization metrics, Δ and STOP states.

1–3 variants, ranked by probability

Earlier drafts fixed the output at exactly three scenarios (optimal / median / minimal). That form is superseded: "always three" presumes three justified variants always exist. Where the data supports only one, the other two would have to be fabricated — violating Invariant 2.

Canonical rule. IAM outputs **from 1 to 3 variants**, always in descending order of calculated probability — the most probable first. Each variant carries:

- **Probability** — the calculated figure, not a verbal confidence.
- **Grounds** — which data and which logic lead to this variant.
- **Limits** — the conditions under which the variant would be wrong.
- **Error margin** — how uncertain the estimate itself is.
- **Risk block** — threat map for that specific state.

Adaptive odd-numbered evaluation

The number of evaluation elements is not fixed in advance. The query determines the minimum analytical resolution required; IAM then brings that count to an **odd** number before evaluating. Odd counts remove symmetric deadlock — an even split can hang the model with no basis to prefer any variant. Oddness guarantees at least a minimal margin — but the margin's *strength* is also measured: a 3/5 majority is minimally sufficient, 5/5 is maximal. IAM reports not only the winner, but the strength of preference, the probability, the error and the data quality.

Layered narrowing

Variants are not all processed to the same depth. A broad primary pass discards the clearly weak or irrelevant; the surviving, narrower set then goes through deeper verification. Lower layers give breadth, upper layers give depth. The final answer is not the whole analytical field — it is what survived successive narrowing.

Gateway Pipeline



no data enters the core bypassing this chain

No data reaches the IAM core bypassing this chain.

- **Ingest Gateway.** Connectors to streams; normalization of events (time, source, type, object).
- **SEVR / AEAD Gateway.** Semantic validation, cross-source consistency, trust level, anomaly detection. Explicit anomalies are rejected.
- **SSR Gateway (data quality).** Below threshold the pipeline halts and returns a request for more data instead of an analytical conclusion — the right to silence, enforced mechanically.
- **IAM Core (A · B · C).** Processing of verified data; construction of the ranked variants with a risk block for each.
- **Action Gateway.** The single exit to the outside world. Variants with source attribution are passed to the DAO. There is no return channel from DAO into IAM.

Three-core structure

Core A — active

Current system state. Primary real-time analytical processing.

Core B — deferred snapshot

Versioned state, used for comparison and deviation detection.

Core C — isolated recovery archive

A separate server prioritising security. Used only for restoration after critical failure.

The right to stop

IAM has an architecturally fixed right not to answer when available data is insufficient to reach a conclusion at an acceptable level of confidence. Silence is not a fault — it is a designed signal.

Better an honest stop than a beautiful wrong answer.

Silence is never mute. A stop must state:

- the reason for stopping;
- exactly which data is missing;
- whether an answer is possible at the current level of incompleteness;
- the risk of error if answered without the missing data;
- a request to the operator or DAO for the next step.

Consequences. The corresponding operation halts. The protocol allows `3600 s` (1 hour) to supply complete data. If it is not supplied, the matter escalates to the DAO.

Seven architectural invariants

These constraints are immutable. No configuration, no LLM, no DAO decision can bypass them. Each is enforced by an automated test that blocks deployment on violation.

#	INVARIANT	MEANING
1	Sovereignty of action	Any action that changes system state is executed only after DAO approval — except deterministic protocol rules, where IAM executes a prescribed rule, not a judgement.
2	Zero tolerance for fabrication	On data deficit the result is “unknown” plus the list of data required. Missing facts are never modelled as if present.
3	Verifiability	Every statement that affects a decision is bound to its sources and is reproducible.
4	Engine tolerance	LLMs and other analytical tools are replaceable. The invariants are not.
5	Audit & provenance	Critical events carry provenance: what, when, why, on which data.
6	Real-time re-evaluation	On new data IAM re-evaluates variants without anchoring to its previous conclusion.
7	Independence from EPR	EPR defines the rules of the meta-network — not the analytical methodology of IAM.

IAM and DAO

The link is one-directional: IAM → DAO. The DAO has zero influence on IAM's analytical methodology. It cannot "order" a preferred variant. If IAM is silent, the DAO receives a data-deficit signal and the escalation that follows it.