BILYTICA AI OPERATING SYSTEM

AI Value Framework

From use-case long list to booked benefit, with unit economics attached at every step.

VALUE · Decide

WHAT IT IS FOR

AI budgets die when nobody can say what a workload costs per transaction or who books the benefit. VALUE sizes, sequences and then tracks realisation against the same numbers used to approve it.

WHO RUNS IT

Owned by the business sponsor with finance as co-signer on the benefit definition.

TYPICAL TIMELINE

Three to four weeks for a portfolio pass; benefit tracking then runs quarterly.

STAGES AND EXIT GATES

1. Discover — Collect candidate workloads from process pain, not from technology lists.
Exit gate: Each candidate names a decision and the person who makes it today.
Artefacts: Use-case long list · Decision owners

2. Size — Estimate benefit, cost to build, cost to run per transaction and the data dependency.
Exit gate: Run cost per transaction is estimated before build is approved.
Artefacts: Value case · Unit economics sheet

3. Sequence — Rank on value, data readiness and risk tier — not on enthusiasm.
Exit gate: Nothing enters build on an uncertified data domain.
Artefacts: Portfolio ranking · Dependency map

4. Commit — Fund a slice with a named benefit owner and an agreed measurement method.
Exit gate: Finance accepts the measurement method before build starts.
Artefacts: Benefit charter · Measurement method

5. Realise — Track actual benefit and actual run cost quarterly, and retire what does not earn.
Exit gate: Every live workload has a current benefit and cost reading.
Artefacts: Realisation report · Retire / scale decision

STANDARDS CROSSWALK

STANDARD	WHAT IT SATISFIES
FinOps for AI	Unit economics, showback and cost anomaly review per workload.
ISO/IEC 42001	Clause 6 objectives and planning, and performance evaluation in clause 9.
NIST AI RMF	Measure function — defining what is measured and how often.

WHAT YOU GET

WHAT WE REFUSE TO DO

- **Sized use-case portfolio** — Every candidate with benefit, build cost, run cost and data dependency.
- **Unit economics model** — Cost per transaction broken down across inference, retrieval and review.
- **Benefit charter** — The measurement method finance agreed to before build.
- **Benefit with no owner** — If no budget line changes, the benefit was never real.
- **Ignoring run cost** — Workloads approved on build cost alone become the next cost-cutting target.
- **Portfolio by enthusiasm** — The loudest sponsor gets sequenced first and the data is not ready.

FREQUENTLY ASKED

How early should run cost be modelled?

Before build approval. Retrieval, review and pipeline costs usually exceed the model call at production volume.

What if the benefit is qualitative?

Then it is stated as a service-level or risk-reduction target with a measurable proxy, and labelled as such.

How do you size AI benefit without inventing numbers?

Sizing uses your own baselines — volumes, cycle times, error and rework rates — and each case is signed by the business owner who will be measured on it.

Does the framework include run cost?

Yes. Inference, retrieval, human review and platform cost are sized alongside benefit, following FinOps-for-AI practice, so a case cannot look positive on build cost alone.

Book the Data Foundation Assessment

Full framework detail: https://bilytica.com/frameworks/ai-value-framework/

bilytica.com
info@bilytica.com