LYBOAI · VALUE-AI FRAMEWORK v1.3

AI Use-Case Assessment Report

Invoice triage assistant — a fictional sample demonstrating the v1.3 consulting-grade outputs: the Assurance Wheel, the weakest-link decision matrix, verdict banners and the Oversight Ladder.

Sample data — fictional organisation

Generated by the VALUE-AI Platform

SAMPLE

1 · EXECUTIVE SUMMARY

Invoice triage assistant

Automated classification and routing of supplier invoices with human review before posting. Assessed on the balanced weighting profile; 16 dimensions answered.

VALUE-AI SCORE	RISK TIER	RISK OVERLAY	WEAKEST LINK
74 /100	Medium	● Amber	L1
Promising	separate axis — risk can veto value	risk needs active management	Risk & responsible AI · 55/100

DEPLOYMENT RECOMMENDATION

Discovery

Close the weakest pillars — Lifecycle Controls first — then run a scoped Proof of Value. A strong business case does not rescue unproven controls: the weakest link, set against what is at stake, drives the decision.

Why the recommendation is capped

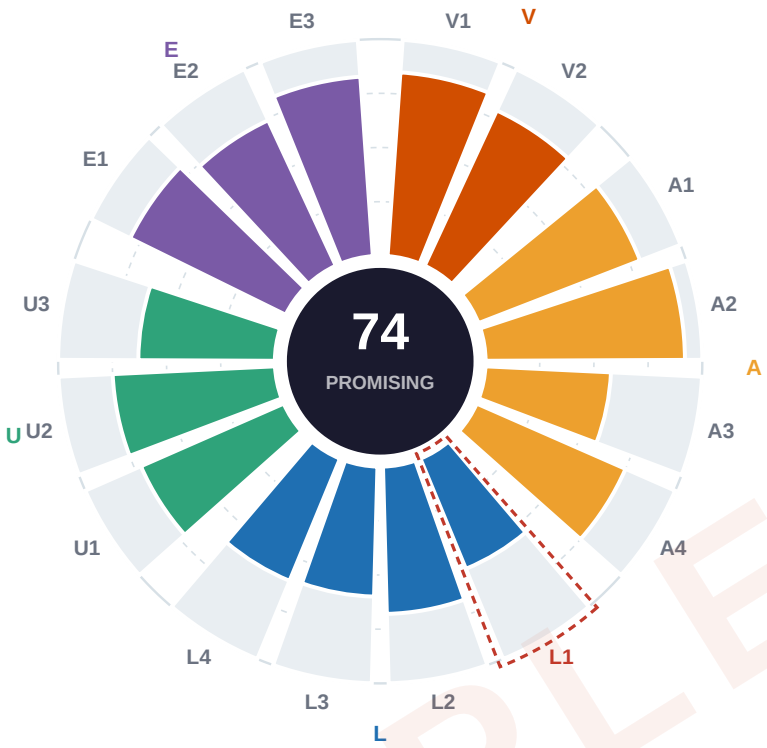
- Lifecycle Controls at 61 on a Medium risk tier — governance, monitoring and the human-review model need definition before pilot.
- Data readiness (A3, 58): currency and representativeness of supplier data are unproven.
- Accountability risk class rated Medium residual with Weak controls — resolve ownership before any deployment gate.

SAMPLE

2 · THE ASSURANCE WHEEL

One picture of readiness

Sixteen dimensions across the five pillars, each spoke filled to its score. Rings mark 25 / 50 / 75 / 100. The dashed outline marks the weakest link — the dimension that caps the recommendation.



Pillar scores

V · Value		82
A · AI Fit		76
L · Lifecycle Controls		61
U · User Readiness		70
E · Evidence		79

Weighted equally across the five pillars (balanced profile). Minimum pillar thresholds apply regardless of the overall score.

SAMPLE

3 · THE DEPLOYMENT DECISION

The weakest-link rule

Risk tier × the weakest pillar drives the recommendation — strength on four pillars never averages away weakness on the fifth. This use case sits at the outlined cell: Lifecycle Controls (61) on a Medium tier.

RISK TIER ↓ · WEAKEST PILLAR →	< 40	40–59	60–79	80+
Low	Discovery	Discovery	Proceed	Proceed
Medium	Govern first	Discovery	Conditions ◉	Proceed
High	Stop	Govern first	Discovery	Conditions
Restricted	Stop	Stop	Govern first	Govern first

Indicative defaults from the published VALUE-AI thresholds and overrides; the live risk overlay (critical risk classes) can cap any cell further.

Critical risk classes (separate axis)

RISK CLASS	RESIDUAL	CONTROL STRENGTH	STATUS
Privacy	Medium	Adequate	Watch
Safety	Low	Strong	OK
Legal / Regulatory	Medium	Adequate	Watch
Accountability	Medium	Weak	Fix first

Any class at High residual with Weak/None controls triggers the hard-stop override — risk can veto value.

SAMPLE

4 · DIMENSION DETAIL

Where the score comes from

DIMENSION	PILLAR	SCORE	PROFILE
V1 · Business value	V	85	
V2 · Strategic alignment	V	79	
A1 · AI suitability	A	80	
A2 · Alternative-solution test	A	92	
A3 · Data readiness	A	58	
A4 · Technology readiness	A	75	
L1 · Risk & responsible AI — weakest link	L	55	
L2 · Privacy & security	L	68	
L3 · Governance & ownership	L	60	
L4 · Human oversight	L	61	
U1 · Process readiness	U	72	
U2 · User readiness	U	75	
U3 · Change & adoption readiness	U	63	
E1 · Proof of Value readiness	E	80	
E2 · Measurement & benefits	E	74	
E3 · Scale readiness	E	83	

SAMPLE

5 · OVERSIGHT & NEXT STEPS

The Oversight Ladder (v1.3)

How much the system may decide on its own, and the approval each level (OL1–OL5) requires. Adapted from the public-domain NIST AI RMF (GOVERN 3.2, MAP 3.5, GOVERN 2.3). This use case operates at **OL3 · Review** — every consequential output passes a person before posting.

OL1	**Assist** AI informs; a person makes every decision.	approval: **Use-case owner**
OL2	**Recommend** AI proposes; a person approves each output before it has effect.	approval: **Use-case owner**
OL3	**Review** AI produces the outcome; a person reviews every consequential output before impact.	approval: **Executive sponsor**
OL4	**Sample** Outputs take effect; people sample and audit after the fact.	approval: **Executive + governance gate**
OL5	**Autonomous** No routine human review of individual outputs — low tier, reversible outcomes only.	approval: **Executive + governance gate, re-approved each cycle**

Recommended next step

Run a four-week Discovery: define the human-review model and audit trail (L1, L3, L4), verify supplier-data currency and representativeness (A3), name the accountable owner with a mandate — then re-assess and design a controlled Proof of Value with a success metric, baseline and counter-metric.

About this report

Generated by the VALUE-AI Platform using Framework v1.3 — evidence-anchored questions (never self-ratings), a non-linear point scale, minimum pillar thresholds and a separate risk axis with hard-stop overrides. VALUE-AI is an independent method adapted from the public-domain NIST AI RMF; it is not certified by any standards body and does not guarantee compliance with any law. All data in this sample is fictional.

SAMPLE