VALUE-AI Framework v1.4

Quick reference · Decide whether an AI idea is worth building — and what must be true before it scales

FIVE PILLARS · SIXTEEN DIMENSIONS

V Value

Is there a real, owned, measurable business problem worth solving?

V1 Business value
Frequency and quantified cost of the problem, and who owns and is asking for it

V2 Strategic alignment
Link to stated strategy, fit with other initiatives, leadership appetite

A AI Fit

Is AI genuinely the right solution — and is it feasible with the data and systems we have?

A1 AI suitability
Whether the task class genuinely suits AI, and whether there is proven precedent

A2 Alternative-solution test
Whether rules, automation, integration, search or process redesign would solve it better

A3 Data readiness
Whether the data exists, is current, and is representative

A4 Technology readiness
Integration effort, hosting constraints, and who supports it in production

L Lifecycle Controls

Can we run it safely, lawfully, and accountably — before, during and after deployment?

L1 Risk & responsible AI
Worst realistic failure, error catchability, explainability need

L2 Privacy & security
Data sensitivity, lawful basis, access control

L3 Governance & ownership
Named owner, approval gate, monitoring and auditability

L4 Human oversight
A defined in-the-loop model with override and escalation

U User Readiness

Are the people, workflows and change plan ready for it to actually be used?

U1 Process readiness
Whether the target workflow is mapped, stable, and has an insertion point

U2 User readiness
Change capacity, and whether users understand limits and their accountability

U3 Change & adoption readiness
Enablement, rollout, communication and support planning

E Evidence

Can we prove value with a controlled Proof of Value before we scale?

E1 Proof of Value readiness
Time to a result, dependencies, and scope discipline

E2 Measurement & benefits
Baseline, agreed metric with target and owner, evaluation method

E3 Scale readiness
Whether PoV success extends to production volume, design and cost

HOW IT SCORES

You answer factual questions about your situation — never self-ratings. Each answer maps to 1, 2, 4 or 5 points (worst to best), rolled up to dimension, pillar and an overall score out of 100. Two things stop a good average hiding a real problem: **the weakest link** (a pillar below 40 holds back a Proceed) and **the risk override** (high residual risk with weak or no controls caps the verdict at Govern first, whatever the score).

Strong · 80+
Proceed — design the Proof of Value now

Promising · 60+
Discovery — close the weakest pillars, then run a scoped PoV

Not ready · 40+
Redesign / prepare — fix data, process or controls first

Unsuitable · under 40
Stop — do not proceed as an AI use case now

VERDICTS

Proceed — design the Proof of Value now

Discovery — close the weakest pillars, then run a scoped PoV

Redesign — fix data, process or controls first

Govern first — resolve controls and risk before any pilot

Consider non-AI — rules, automation or a process fix may do it better

Stop — not an AI use case now

USE IT IN FOUR STEPS

1 · Pick one use case. Something real — built, bought or an agent.

2 · Answer as things are today, not as you hope they'll be.

3 · Read the weakest link first. That's the conversation to have.

4 · Agree stop criteria before any pilot starts (E1.5).

VALUE-AI Framework v1.4

Quick reference · Decide whether an AI idea is worth building — and what must be true before it scales

OVERSIGHT LADDER — HOW MUCH MAY IT DECIDE ON ITS OWN?

Level	Name	What it means	Who approves	What to expect
OL1	**Assist**	AI informs; a person makes every decision.	Use-case owner	Basic usage guidance and quality sampling.
OL2	**Recommend**	AI proposes; a person approves each output before it has effect.	Use-case owner	Reviewer training and a workload check — approval must stay meaningful at volume.
OL3	**Review**	AI produces the outcome; a person reviews every consequential output before impact.	Executive sponsor	A defined review point, tested override, and a named person authorised to pull it.
OL4	**Sample**	Outputs take effect; people sample and audit after the fact.	Executive sponsor + governance gate	Stated sampling rates, drift monitoring, an incident route, and a tested kill-switch.
OL5	**Autonomous**	No routine human review of individual outputs.	Executive sponsor + governance gate, re-approved each review cycle	Low risk tier and reversible outcomes only, with continuous monitoring — never for decisions with legal or similarly significant effect on individuals.

NEW IN V1.4 — TWELVE QUESTIONS (76 IN THE BANK; EACH ASSESSMENT ASKS ABOUT 40)

A4.6 If it will act in other systems, is the plan to limit its access to only what the task needs?

L1.5 If it will take actions, could those actions be undone?

L1.6 Is there the capability to red-team test it — deliberately trying to break or misuse it — so abnormal behaviour is detected early?

L2.5 If the AI helps make decisions about people, are you ready to tell them so?

L2.6 Will the solution be designed to resist instructions hidden in what it reads (prompt injection) and poisoned data — and do you or your vendor have the capability to build that in?

L3.5 Will every action it takes be traceable to a responsible person — and can your systems support that?

L3.6 Do your AI vendor contracts cover how your data is used, where it is kept, and who is liable?

L3.7 Do you know which AI laws and rules apply where you operate or sell?

L4.5 Will it have hard limits — on spending, volume or the kinds of action it can take — with points where a person must approve, and a way to stop it?

U2.4 Will people know when they are dealing with AI, or reading AI-generated content?

U3.4 Do you know which AI tools staff already use, and have they been trained to use AI safely?

E1.5 Have you agreed in advance what result means scale, change or stop — and by when?

For AI that acts on its own, access, reversibility, hidden-instruction protection and hard limits (A4.6, L1.5, L2.6, L4.5) are always asked. An agent that can take irreversible actions without limits or injection protection is capped at **Govern first**.

DRAWS ON — FREE TO READ

NIST AI RMF — The Govern–Map–Measure–Manage structure for identifying, measuring and managing AI risk.

VALUE-AI Method — LyboAI's plain-language operating disciplines — original wording and structure, adapted from the public-domain NIST AI RMF.

Australia's Guidance for AI Adoption — The National AI Centre's six essential practices ('AI6', Oct 2025, building on the earlier Voluntary AI Safety Standard): decide who is accountable, understand impacts, measure and manage risks, share essential information, test and monitor, maintain human control.

OECD AI Principles — Human-centred, trustworthy AI values: accountability, transparency, fairness and safety.

Australia's AI Ethics Principles — The eight voluntary Australian principles — wellbeing, human-centred values, fairness, privacy, reliability & safety, transparency & explainability, contestability, accountability.

Singapore Model AI Governance Framework for Agentic AI — The first government framework for AI agents (Jan 2026): bound an agent's risks up front, keep humans meaningfully accountable, apply technical controls across its lifecycle, and enable end-user responsibility.

OWASP Top 10 for LLM & Agentic Applications — The open security community's threat lists for AI apps and agents — prompt injection, goal hijack, tool misuse, identity and privilege abuse, memory poisoning, cascading failures.

Lean validation — Build–Measure–Learn validation with explicit decision gates before scaling.

DATES TO HAVE ON THE CALENDAR

10 Dec 2026 — Australian privacy policies must describe automated decisions that significantly affect people. **2 Dec 2027** — EU AI Act obligations for high-risk (Annex III) systems begin, under the 2026 simplification package. Check official sources; dates move.