AI Assurance Design Checklist

Ten questions to answer before an AI system informs a consequential decision. Use it before a pilot moves to production, or whenever a regulator, auditor, or board asks how an AI-supported decision was reached.

- ☐ 01 **Decision inventory.** Do you know every decision this system informs, who owns each one, and what happens if it is wrong?
- ☐ 02 **Risk tier.** Has each decision been classified by its impact on customers, finances, legal exposure, and safety, with controls proportionate to that tier?
- ☐ 03 **Requirements up front.** Are explainability, traceability, and human oversight written as system requirements, not as policies to add after deployment?
- ☐ 04 **Traceability.** For any output, can you reconstruct the inputs, model version, configuration, and data sources that produced it?
- ☐ 05 **Human oversight.** Is it defined which decisions require human review, who reviews them, and what triggers an escalation?
- ☐ 06 **Data governance.** Are data sources, permitted uses, and retention of inputs and outputs documented and enforced?
- ☐ 07 **Realistic testing.** Has the system been evaluated on decision scenarios from your own domain, including edge cases and known failure modes?
- ☐ 08 **Vendor contracts.** Do your agreements match how the technology works: data flows, subprocessors, model changes, audit rights, and exit?
- ☐ 09 **Monitoring.** Are drift, error rates, and human overrides measured, with a named owner and a regular review cadence?
- ☐ 10 **Accountability.** If a regulator asked tomorrow, could you explain how one specific decision was reached and show the record?

If any answer is "not yet," that is where to start. For a design review or an executive session on AI assurance, book a complimentary introductory call at technology-operation.com or write to valid.0.western@icloud.com.

© 2026 Federico J. Quijada · www.technology-operation.com