BILYTICA AI OPERATING SYSTEM

AI Assurance Framework

One control set that answers NIST, ISO 42001, the EU AI Act and the Saudi regulators at once.

TRUST · Assure

WHAT IT IS FOR

Assurance fails when each regulator triggers a separate project. TRUST defines controls once, then maps the same evidence to every framework a reviewer might ask about.

WHO RUNS IT

Owned by risk and compliance, operated with the delivery team that builds each workload.

TYPICAL TIMELINE

Control set stood up in four to six weeks; evidence then accrues automatically per release.

STAGES AND EXIT GATES

1. Classify — Assign a risk tier per workload from impact, autonomy and data sensitivity.

Exit gate: The tier is recorded before any model touches production data.

Artefacts: Risk tier record · Impact assessment

2. Control — Apply the control set that the tier requires — no more, no less.

Exit gate: Every required control has an owner and a test.

Artefacts: Control matrix · Control tests

3. Evidence — Capture logs, evaluations, approvals and data lineage as the release happens.

Exit gate: Evidence is produced by the pipeline, not assembled before an audit.

Artefacts: Evidence pack · Model card

4. Attest — Sign off per release against the mapped clauses, with named approvers.

Exit gate: No production release without a current attestation.

Artefacts: Attestation record · Clause crosswalk

5. Monitor — Watch drift, incidents and control failures, and re-tier when behaviour changes.

Exit gate: A failed control raises an incident with a defined response time.

Artefacts: Drift and incident log · Re-tier decision

STANDARDS CROSSWALK

STANDARD	WHAT IT SATISFIES
NIST AI RMF	Govern, Map, Measure and Manage functions map to Classify, Control, Evidence and Monitor.
ISO/IEC 42001	AI management system clauses 4-10 and Annex A controls.
EU AI Act	Risk tiering, technical documentation and post-market monitoring duties.
SDAIA AI Ethics	Fairness, transparency, accountability and human oversight principles.
PDPL / NDMO	Personal data handling, minimisation and cross-border rules per workload.
NCA ECC	Identity, logging, key management and change control for AI systems.

WHAT YOU GET

- **Tiered control matrix** — Controls per risk tier, each with an owner and a test.
- **Audit evidence pack** — Release-time evidence mapped to every referenced clause.
- **Standards crosswalk** — One table showing which control answers which framework.

WHAT WE REFUSE TO DO

- **Audit-time archaeology** — Reconstructing evidence after the fact costs more than generating it during release.
- **One control set for everything** — Applying high-risk controls to every workload stalls low-risk work and breeds workarounds.
- **Policy without tests** — A control that is not tested is a sentence, not a safeguard.

FREQUENTLY ASKED

Do we need ISO 42001 certification to use this?

No. The framework is structured so the same controls and evidence would support a certification effort if you later choose one.

How are risk tiers decided?

From three inputs: impact on a person or asset, the degree of autonomy granted, and the sensitivity of the data involved.

How long does it take to stand up the control set?

Four to six weeks for the control set itself; after that, evidence accrues automatically with each release rather than being reconstructed at audit time.

Does one control set apply to every AI system?

No. The Classify stage assigns a risk class first, and controls scale with that class, so a low-risk assistant is not governed like a decisioning system.

Book the Data Foundation Assessment

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

bilytica.com
info@bilytica.com