

AI GOVERNANCE HANDBOOK

YOUR AI PRE-FLIGHT CHECKLIST

A guide to governing AI use cases, from the register through to continual improvement. The principles borrowed from aviation is simple - accountability never transfers to the machine. Additionally, Humans-in-the-Loop is important but only meaningful when we have the authority, capability and evidence to challenge what AI is telling us. Most importantly, keep flying the plane when it all goes wrong.

infotrust

SECURE YOUR FUTURE.

THE AI GOVERNANCE CYCLE

Starting Point

Your AI use cases and governance framework

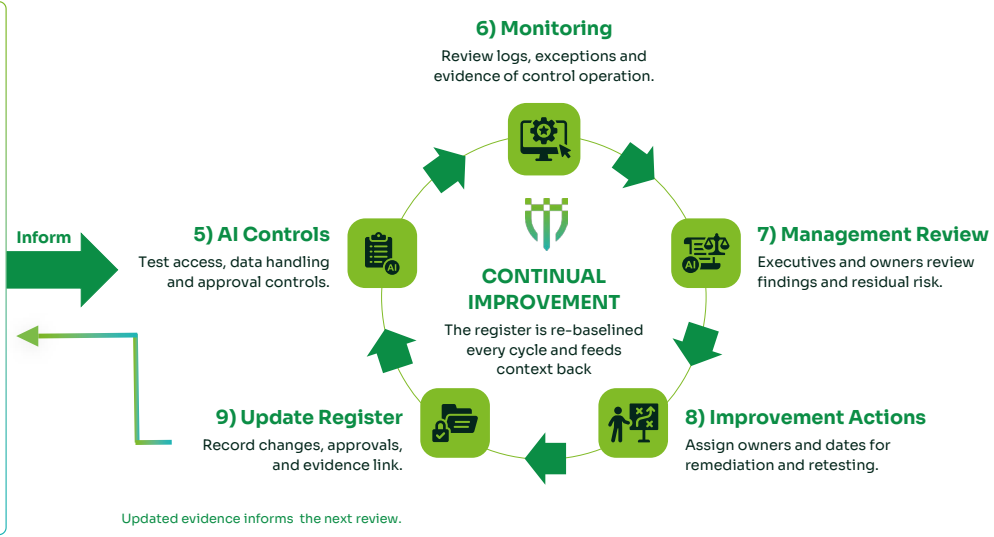
1) AI Resources Register
Record approved tools and the applications they connect to.

2) AI Use Case
Confirm purpose, owner, and data for each AI use case.

3) AI Risk Assessment
Prioritise financial, data, and operational impacts.

4) AI Impact Assessment
Check decision impacts and human review before release or change.

Inform



Step	What Good Looks Like	Evidence to Keep (Your AI Flight Recorder)
1. AI Resources Register	Record approved AI tools, agents and the applications, data, people and dependencies they connect to. A use case should not hide behind a "productivity tool" label.	Register entries linking software, infrastructure, people and dependencies to use cases, risks and controls.
2. AI Use Case	Describe each use case in two plain-English sentences: purpose, owner and data. Name who can pause or stop it, including vendors or MSPs.	Use case statement, business owner, risk owner, stop and override authority.
3. AI Risk Assessment	Prioritise financial, data and operational impacts against your risk appetite.	Risk rating, tolerance and acceptable limits.
4. AI Impact Assessment	Check decision impacts and human review before release or change. Reviewers need the domain expertise to challenge outputs.	Impact assessment, review sign-off and a manual fallback plan.
5. AI Controls	Test access, data handling and approval controls. Set performance measures and escalation triggers.	Model or service version, permissions, control test results and triggers.
6. Monitoring	Review logs, exceptions, overrides and evidence that controls operate.	Outputs, checks, warnings, overrides, handoffs and escalations.
7. Management Review	Executives and owners review findings and residual risk.	Review records and residual risk decisions.
8. Improvement Actions	Assign owners and dates for remediation and retesting. Confirm completion and effect, not just acknowledgement.	Action log with owner, deadline and confirmed result.
9. Update Register	Record changes, approvals and evidence links. Re-baseline every cycle.	Updated register with approval and evidence links.

Updated evidence informs the next cycle. Revisit assumptions, acceptable limits and controls whenever the use case, operating conditions or results change. Traceability test: can you follow the source evidence to human approval, then to the instruction given to the action owner, and finally to the confirmed result? If not, your flight recorder has gaps.

When AI Gets it Wrong:

Aviate, Navigate, Communicate

AVIATE

Retain Control

Stabilise operations and protect critical services and evidence. Know who can stop or override AI actions. Keep your incident response, business continuity and disaster recovery plans connected.

NAVIGATE

Exercise Judgement

Challenge AI outputs against authoritative sources and the context of use. Check what data was used, what remains uncertain and what contradicts it.

COMMUNICATE

Explain Evidence & Uncertainty

Label each statement as verified fact, AI assessment or untested assumption. Name the decision-maker, action owner and deadline. Confirm the work is complete and effective.

Pre-flight Check:

Five Questions for One AI Use Case

1. Do we know the AI system, its uses and the data it can access, and is the use case visible in the register?
2. Who owns the system, the risks and the decisions, and who can actually stop the AI?
3. What evidence verifies the output? Is the source current, and what evidence would exist if the control were not working?
4. Can we spot errors and work without AI? Exercise it with a plausible wrong output, for example in a tabletop exercise.
5. Do our records explain decisions and overrides? Keep the original output, source evidence, version, reviewer, escalation, instruction and confirmed result.

Aligned with the intent of ISO/IEC 42001 (AI management systems), the NIST AI Risk Management Framework and Australian Signals Directorate guidance on engaging with AI. This handbook helps you connect the register to risk appetite, impact assessment, controls, monitoring, management review and corrective action. Remember that effective preparedness happens on the ground, before take-off, not on the fly.

For inquiries, contact us today:

info@infotrust.com.au infotrust.com.au

infotrust