

Principal Brief & Risk Checklist

Governance Framework for AI in ASIC-Regulated Firms

Executive Summary

This document outlines the governance controls, risk boundaries, and implementation protocols required to deploy AI drafting support safely within an Australian Financial Services Licence (AFSL) environment.

Prepared for Principals & Responsible Managers

December 17, 2025

1 Risk Framework: The "Safe Harbor" Approach

Deploying AI in financial advice requires a strict separation between **administrative drafting** (safe for AI) and **advice judgement** (human only). BackPro is architected to enforce this separation.

1.1 What Doesn't Change

Regardless of the technology used, the following core principles remain immutable:

- **Best Interests Duty (BID):** Remains 100% with the adviser. AI cannot hold a licence or fiduciary duty.
- **Accountability:** The signing adviser and AFSL are responsible for every word in the SOA, regardless of who (or what) drafted it.
- **Human Review:** No document can be sent to a client without authorised human review and sign-off.

1.2 Where AI Helps Safely

We restrict AI usage to low-risk, high-volume tasks:

- **Drafting Standard Text:** Generating standard strategy explanations using firm-approved templates.
- **Policy Retrieval:** Summarising internal compliance manuals and product lists.
- **Consistency Checking:** Identifying discrepancies between Fact Find and SOA.

Explicit "Red Lines" (What We Do Not Do)

- **NO Autonomous Advice:** The AI never recommends products or strategies independently.
- **NO Public Data Training:** Client data is never used to train public models (e.g., ChatGPT).
- **NO "Black Box" Logic:** All AI outputs must be traceable to source documents or templates.

2 Compliance & Governance Checklist

Use this checklist to validate readiness for AI implementation within your practice.

1. Data Sovereignty & Security

- Data Location:** All client data remains within Australian jurisdiction (or approved equivalent).
- Encryption:** Data is encrypted at rest and in transit.
- Isolation:** Your firm's data is isolated from other tenants (no cross-contamination).
- Access Control:** MFA and role-based access control (RBAC) are enforced.

2. Audit Trail & Oversight

- Version History:** System logs exactly what the AI generated vs. what the human edited.
- Source Attribution:** AI responses link back to the specific policy or template used.
- Review Logs:** Timestamps of when the adviser reviewed and approved the content.

3. Operational Governance

- Staff Training:** All users have completed "AI Boundaries" training (what it can/cannot do).
- Template Approval:** Only Compliance-approved templates are loaded into the system.
- Escalation Path:** Clear process for reporting AI hallucinations or errors.

3 Due Diligence Questions for Vendors

When evaluating any AI provider, Principals should ask:

1. **Do you train your models on my client data?**

BackPro Answer: No. Your data is used only for retrieval (RAG) within your instance.

2. **Can you produce an audit trail of the AI's contribution?**

BackPro Answer: Yes. We track the "diff" between AI draft and final version.

3. **What happens if the AI makes a mistake?**

BackPro Answer: Our UI forces human review. Liability remains with the adviser, supported by our indemnity framework for software errors.

Next Steps

Ready to validate your governance framework?

Book a 20-minute Compliance Walkthrough to see these controls in action.

[Book Compliance Walkthrough →](#)