AI GOVERNANCE REVIEW PACK · A0

# AI governance review pack: cover and decisions

**Fictional organisation, illustrative only, not a client case study.** Quillmarrow Supplies Pte. Ltd. (fictional) is invented. All people are shown by role only.

| | |
|---|---|
| **Organisation** | Quillmarrow Supplies Pte. Ltd. (fictional) |
| **Document** | A0 · AI governance review pack: cover and decisions |
| **Version** | 1.0 |
| **Date** | 30 October 2026 (illustrative) |
| **Prepared by** | Alex Lew, reviewer and sign-off |
| **Client approver (role)** | Chief Operating Officer |
| **Label** | Fictional organisation, illustrative only, not a client case study |

**Quillmarrow Supplies Pte. Ltd. (fictional)**, an office-supplies distributor and online store with about 80 staff, a Singapore head office and a Malaysia fulfilment partner.

AI and data governance from someone who builds the controls, not just the policies.

## Why this review was commissioned

Staff were using ChatGPT and Microsoft 365 Copilot without clear rules. The helpdesk vendor switched on an AI reply-drafting feature. Purchasing was piloting an agent that can email purchase orders to suppliers. HR proposed a CV-screening service. A large enterprise customer sent a supplier AI questionnaire, and the company had just appointed its own Data Protection Officer. Management wanted to know what its AI may do, and who approves it.

## Decisions this supports

- Which AI uses to approve, restrict or pause.
- What information staff may put into each tool.
- When a person must check an output or approve an action.
- What an agent may read, change or send, and what it must never do.
- Which supplier questions need answers before you proceed.
- Who owns each risk, control and follow-up action.

### Where each decision is answered

| Decision | Answered in | Decided by (role) |
|---|---|---|
| Which AI uses to approve, restrict or pause | A2 register, A3 tiering, A5 impact assessments, A9 decisions needed | COO; MD for Tier 3 uses |
| What information staff may put into each tool | A4 AI-use policy, section 4 table | MD approves the policy |

| Decision | Answered in | Decided by (role) |
|---|---|---|
| When a person must check an output or approve an action | A6 oversight matrix | COO |
| What an agent may read, change or send | A7 specification, AIA-02 | MD, before any sending is enabled |
| Which supplier questions need answers first | A8 questionnaire and its follow-up actions | COO with the use owner |
| Who owns each risk, control and follow-up action | A2 owner roles, A9 30/60/90-day plan | COO |

## The pack at a glance

| Measure | What the pack records |
|---|---|
| AI uses in scope | 20, from catalogue drafting to a supplier reorder agent |
| Working risk tiers | 6 Tier 1 Low, 10 Tier 2 Medium, 4 Tier 3 High |
| Uses not simply approved | 7: two restricted, one on hold, two pilots under conditions, one not yet live, and one not permitted for automatic decisions |
| Impact assessments | 3: helpdesk drafting (AIA-01), reorder agent (AIA-02), CV screening (AIA-03) |
| Evidence status across the register | 4 complete, 14 partial, 2 missing |
| Vendor questions | 30 questions to 3 fictional vendors, with every gap left visible |
| Management decisions requested | 5, each with options, a recommendation, an owner and a deadline (A9) |

These are counts of what the sample documents contain. They are not results, savings or outcomes.

## What is in this pack

| Ref | Deliverable | Format | Approver role |
|---|---|---|---|
| A1 | Engagement scope sheet | PDF, DOCX | Chief Operating Officer |
| A2 | AI use register (20 uses) | XLSX, PDF | Chief Operating Officer |
| A3 | Practical risk tiering method | PDF, DOCX | Chief Operating Officer |
| A4 | AI-use policy | PDF, DOCX | Managing Director |
| A5 | AI impact assessments AIA-01, AIA-02, AIA-03 | PDF, DOCX | COO; Managing Director for Tier 3 |
| A6 | Human oversight decision matrix | XLSX, PDF | Chief Operating Officer |
| A7 | Agent governance checklist: permissions and audit-trail specification | PDF, DOCX | Chief Operating Officer |
| A8 | Vendor and model due diligence: 30 questions, 3 vendors | XLSX, PDF | Chief Operating Officer |
| A9 | Management report and action plan | PDF, DOCX | Managing Director |
| A10 | Optional ISO/IEC 42001 readiness check (excerpt) | PDF, DOCX | Chief Operating Officer |

## What this pack shows

- **One register, owned.** All 20 AI uses in one place, each with a named owner role, a working risk tier with its rationale, the human check that applies, the evidence still missing and the event that forces a re-review (A2, A3).
- **Rules staff can follow.** What information may go into which tool, what is restricted or not permitted, and how to report a problem (A4).
- **People in control of agents.** For the supplier reorder agent, which can email purchase orders, the pack sets who approves before it acts, what is monitored while it runs, the stop conditions and who can halt it, and the audit trail that proves each step (A6, A7, AIA-02).
- **Missing evidence left visible.** Vendor answers are recorded as Evidence provided, Partial or Not provided, and gaps become dated follow-up actions rather than assumptions (A8, A9).
- **Decisions, not just findings.** Five findings, decisions needed with a recommendation and owner, and a 30/60/90-day plan where each action names its closure evidence (A9).

## How to read this pack

1. Start with the **management report (A9)**: the executive summary, the five findings, the decisions needed and the 30/60/90-day plan.
2. Check the **register (A2)** for any use you care about, then the **tiering method (A3)** for why it sits in its tier.
3. For the agent that can email purchase orders, read **AIA-02 (A5)** with the **oversight matrix (A6)** and the **permissions and audit-trail specification (A7)**.
4. For supplier gaps, read the **vendor questionnaire (A8)**.
5. Give staff the **AI-use policy (A4)** once the Managing Director approves it.

## Approvals and status

| Item | Client approver (role) | Status in this sample |
|---|---|---|
| A1 scope sheet | Chief Operating Officer | Agreed before fieldwork |
| A2, A3, A6, A7, A8, A10 | Chief Operating Officer | Issued for approval |
| A4 AI-use policy | Managing Director | Issued for approval; staff rollout after approval |
| A5 impact assessments | COO; Managing Director for Tier 3 (AIA-02, AIA-03) | Issued for decision |
| A9 management report | Managing Director | Issued for decision; decisions due as dated in A9 |

The reviewer signs off each deliverable before issue. Quillmarrow's management makes and records the decisions.

## After the review

- **Owner.** The COO owns the register and the action plan; each use has an owner role in A2.
- **Re-review triggers.** A new tool or vendor feature, a change in permissions, a new data category, an incident, or a vendor change of terms puts a use back through A3.
- **Cadence.** Tier 3 uses are reviewed quarterly, Tier 2 every six months and Tier 1 annually.

- **What is not included.** Legal advice, certification or a compliance guarantee, AI Verify technical testing, penetration testing or security engineering, and changes to live systems, unless separately scoped. No healthcare-sector work or on-site delivery. The full list is at the end of every document.

## What Quillmarrow supplied

- Role holders for interviews: the COO, the Head of Operations (the appointed DPO), the Head of Customer Service, the Purchasing Lead, the HR Manager, the Head of Marketing, the IT Manager and the Finance Director.
- System and licence lists, the helpdesk and reorder-agent settings screens (client-reported), vendor documents received to date, and the enterprise customer's supplier AI questionnaire.
- Workflow walk-throughs using descriptions and made-up or de-identified samples. No customer, staff or supplier records were provided or needed.

## Limits of this pack

- Settings and usage are client-reported unless an evidence item is marked Complete in A2 or Evidence provided in A8.
- No AI Verify technical tests, model testing or security testing were run; the mapping covers principles and process checks only.
- Legal questions, such as vendor contract terms or the basis for using personal data, are flagged for a Singapore-qualified lawyer and are not answered here.
- Missing evidence is listed, owned and dated in A8 and the A9 action plan rather than assumed.

## Reference frameworks used

- Model AI Governance Framework, Second Edition (PDPC and IMDA, January 2020).
- Model AI Governance Framework for Generative AI (IMDA and AI Verify Foundation, May 2024).
- PDPC Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems (March 2024).
- Model AI Governance Framework for Agentic AI (IMDA, January 2026), a voluntary framework.
- AI Verify testing framework, 11 principles. The pack maps to the principles and process checks only; no AI Verify technical tests were run.
- ISO/IEC 42001:2023, used only for the optional readiness excerpt (A10). Readiness is not certification.
- Personal Data Protection Act 2012, named where personal data is involved. Legal questions go to a Singapore-qualified lawyer.

## How the pack was prepared

About two weeks, fully remote, working from descriptions, system lists and made-up or de-identified samples. No client personal data was put into AI tools. Review rounds are a set number agreed in the quote.

Prepared by Alex Lew (CFA, AIGP, CIPP/A, FIP), reviewer and sign-off. I personally review and sign off every deliverable.

Quillmarrow appoints its own Data Protection Officer. You appoint your own DPO. I support them.

## About this sample

**Fictional organisation, illustrative only, not a client case study.** Quillmarrow Supplies Pte. Ltd., its people, vendors, findings and dates are invented. Nothing in this document reports a client result.

I personally review and sign off every deliverable. You appoint your own DPO. I support them.

This is practical compliance support, not legal advice. For legal advice, use a Singapore-qualified lawyer.

**What's not included**

- Legal advice, legal opinions on vendor terms or regulatory representation.
- AI Verify technical testing, model benchmarking, penetration testing or security engineering, unless the relevant work is separately scoped.
- Building an AI system, implementing integrations or changing live permissions and configurations.
- Certification, a compliance guarantee or an assurance that a model cannot fail.
- An organisation-wide ISO/IEC 42001 readiness engagement unless expressly included in the quote.
- Healthcare-sector work or on-site delivery.

ISO/IEC 42001 work is a readiness check only, never certification. Questions about this sample: DM @alewyl1 on X, or ask for a free intro call (coffee chat).