AI GOVERNANCE REVIEW PACK · A3

Practical risk tiering method

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	A3 · Practical risk tiering method
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 uses these tiers to decide how much review and oversight an AI use needs. They are working categories for this business, **not statutory classifications**. A tier does not establish that a use is lawful, that its controls are operating or that it cannot fail.

1. How the tier is set

The use owner and the COO consider the intended workflow, the people and information involved, and the action the system can take. They score each factor from 1 to 3 using the best available evidence. If evidence is missing, they record the gap and use a cautious provisional score until it is resolved.

Factor	1	2	3
Potential harm to people or the business	Limited impact; an error is unlikely to affect an important decision or service.	An error could cause a customer complaint, operational disruption or commercial loss.	An error could materially affect a person, an important decision or a significant business commitment.
Personal data	No personal data intended.	Personal data may be used, but the scope is limited and necessary for the task.	Personal data is central to the use, extensive, or its scope and handling are uncertain.
Autonomy	Suggests or drafts; a person checks before use.	Acts within a bounded workflow, with a person monitoring or reviewing exceptions.	Can take consequential action or materially influence a decision; permissions or boundaries are not established.
Reversibility	An error is easily corrected before it takes effect.	Correction needs effort or follows an external communication.	An action or its effects are difficult to undo once it occurs.

Add the four scores:

Total	Working tier
4–6	Tier 1 Low
7–9	Tier 2 Medium

Total	Working tier
10–12	Tier 3 High

Two rules override the total. A use that makes a **final decision about an individual without human review**, or cannot be stopped, is **Not permitted** until redesigned. An agent that can send, pay or change records is **at least Tier 3 while its permissions are unbounded**.

A high tier is not itself a decision to proceed. The MD approves High-tier uses; the status in the register records whether a use is approved, conditional, a pilot, on hold or restricted. If a calculation and the agreed register tier differ, keep the existing restriction, record the discrepancy and refer it to the COO and, for a High-tier use, the MD. Do not silently lower a tier.

2. What each tier requires

These are minimum requirements for this method. A use owner may add controls where the circumstances justify them. Required controls are not evidence that they already operate.

Tier	Approval to start	Human check	Logging	Review frequency	Impact assessment	Approver
Tier 1 Low	Register the use and document its conditions before use.	Check material outputs before publication or action, or monitor exceptions where the register specifies human-over-the-loop.	Retain the approval and relevant checks or exceptions.	At least annually and on a trigger.	Not normally required; complete one if the risks warrant it.	Use owner.
Tier 2 Medium	Document the purpose, data, permissions and conditions before use.	Apply the register's human review or exception-monitoring arrangement.	Retain approvals, material errors, complaints and interventions.	At least every six months and on a trigger.	Required where the COO identifies material uncertainty or harm; AIA-01 is required for AI-02.	Use owner and COO.
Tier 3 High	Complete the required assessment or oversight design, identify evidence gaps and obtain approval before use or pilot.	Document who approves actions, monitors operation and can stop the use.	Retain decisions, action approvals, exceptions and interventions.	At least every three months, during pilots as agreed, and on a trigger.	Required for AI-03 and AI-04; AI-12 and AI-16 receive oversight design in this engagement. Additional assessment may be required before approval.	MD, with the use owner and COO providing the recommendation.
Not permitted	Do not start or continue the prohibited action.	Maintain a human-controlled alternative where appropriate.	Record the prohibition, any attempted use and the redesign decision.	Reassess only after redesign or a relevant change.	Assess the redesigned use before any proposed approval.	MD decides on any High-tier redesign; the prohibition remains until the relevant decision is made.

3. Worked examples

The scores illustrate this method using the agreed register. They are not evidence that the described controls have been implemented.

AI-01: Catalogue product description drafting

Harm	Personal data	Autonomy	Reversibility	Total	Tier
1	1	1	1	**4**	**Tier 1 Low**

Fictional Draft Studio prepares catalogue copy; no personal data is intended for the task. The Head of Marketing checks factual claims and copy before publication. An incorrect draft can be corrected before it reaches buyers. Prompt, access-setting and review records have not been supplied, so the approval conditions still need evidence.

AI-02: Helpdesk reply drafting

Harm	Personal data	Autonomy	Reversibility	Total	Tier
2	3	1	2	**8**	**Tier 2 Medium**

Customer conversations may contain personal data, and an incorrect reply could mislead a customer. A Customer Service user must check and approve each reply before it is sent; a sent reply can be corrected, but not recalled. AIA-01 is identified. Settings, data flow and evidence of review have not been supplied, so the intended approval gate must not be treated as verified.

AI-03: Supplier reorder agent pilot

Harm	Personal data	Autonomy	Reversibility	Total	Tier
3	2	3	3	**11**	**Tier 3 High**

The agent can prepare or email purchase orders, which may create commitments that are difficult to undo. Supplier contact data may also be involved. The pilot requires the Purchasing Lead's approval before each order is sent. Technical bounds, approval-workflow evidence and test evidence have not been supplied; the autonomy score remains cautious until those controls are demonstrated. AIA-02 is identified.

AI-16: Credit-limit recommendation for B2B accounts

Harm	Personal data	Autonomy	Reversibility	Total	Tier
3	2	2	3	**10**	**Tier 3 High**

A recommendation could materially affect a business customer's access to credit, and an unsuitable decision may be difficult to unwind. The Finance Director or authorised finance staff make the decision. **Automatic decisions are not permitted** for this use. The vendor, inputs, recommendation logic and decision records have not been supplied; the score is provisional and does not imply that those controls have been verified.

4. When to re-tier

The use owner should refer the use back to the COO when its purpose, users, data categories or affected people change; a vendor enables a new feature; the system gains permission to send, pay or change records; or a pilot expands. Re-review is also needed after a complaint, incident, harmful output, failed control or newly

discovered evidence gap. Where the use is or may become High-tier, refer the decision to the MD. Keep existing restrictions in place while the review is pending.

5. How this relates to the references

The **Model AI Governance Framework, Second Edition (PDPC and IMDA, January 2020)** considers the probability and severity of harm when determining whether people should be in, over or out of the decision loop. Quillmarrow's scores are a practical local aid to that discussion, not a replacement for it.

For generative uses, the **Model AI Governance Framework for Generative AI (IMDA and AI Verify Foundation, May 2024)** provides broader considerations, including data, security and accountability. For agents, the **Model AI Governance Framework for Agentic AI (IMDA, January 2026)** emphasises assessing and bounding risks upfront. Mapping a use to these references is not proof that its controls work.

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).