AI GOVERNANCE REVIEW PACK · A4

AI-use policy

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	A4 · AI-use policy
Version	1.0
Date	30 October 2026 (illustrative)
Prepared by	Alex Lew, reviewer and sign-off
Client approver (role)	Managing Director
Label	Fictional organisation, illustrative only, not a client case study

1. Purpose and who this applies to

This policy tells you how to use AI for Quillmarrow work, what information you may enter, what checks you must make and when to ask for help. You remain responsible for your work and for decisions made using an AI output.

It applies to all Quillmarrow staff and contractors, including the outsourced web developer. It also applies to Tanjung Contoh Logistics Sdn. Bhd. when it uses Quillmarrow data in its Malaysia fulfilment work, subject to the agreed partner arrangements.

Use this policy alongside the AI use register and your normal information-security and operational procedures. If instructions conflict or you are unsure whether a use is approved, stop and ask your use owner or the COO. Do not assume that a tool is approved because it is available on your device or offered by a vendor.

This policy is for management approval. The restrictions and approval conditions in the agreed AI use register must remain visible when the policy is implemented. Quillmarrow appoints its own DPO. The Head of Operations is the appointed DPO; the reviewer supports Quillmarrow and is not its DPO.

2. Key terms

Term	Meaning in this policy
AI tool	Software or a feature that produces content, predictions, classifications, recommendations or other outputs using AI. It includes AI features within existing business systems.
Approved tool	A tool or feature that Quillmarrow has authorised for a particular purpose, users and information categories. Approval for one use does not approve every possible use.
Agent	An AI-enabled system that can carry out steps or take actions towards a task, such as preparing or sending a purchase order.
Personal data	Information about an identifiable individual. This can include customer, staff, applicant or supplier-contact information.
Confidential information	Non-public information that could harm Quillmarrow or another party if disclosed, including supplier commercial terms and non-public business plans.

Term	Meaning in this policy
Output check	A meaningful review by a person who can question, correct, reject or stop an AI output before it is used or acted upon.

An output check is more than clicking an approval button. You must have enough information, time and authority to make the check and to reject the result.

3. Approved tools and uses

Approval is for the stated use and conditions only. The use owner can tell you the approved users, information categories and workflow. If settings or supporting evidence remain unconfirmed, do not assume that a technical control is active.

ID	Tool or use	Status and what you must do
AI-01	Catalogue drafting — Fictional Draft Studio	**Approved with conditions.** The Head of Marketing checks factual claims and copy before publication.
AI-02	Helpdesk reply drafting — Fictional Helpdesk	**Approved with conditions.** A Customer Service user checks and approves each reply before it is sent.
AI-03	Supplier reorder agent — Fictional Reorder Agent	**Pilot with approval gate.** The Purchasing Lead approves every purchase order before sending; the agent controls in section 7 apply.
AI-04	CV screening and candidate ranking — Fictional Talent Screen	**On hold; recommended not to proceed in its current form.** Do not use it for screening or ranking while the hold remains.
AI-05	Microsoft 365 Copilot for email, documents and spreadsheets	**Approved with conditions.** Check outputs and recipients before use or sharing; permitted settings and access remain to be confirmed.
AI-06	ChatGPT on staff personal accounts	**Restricted.** Do not put work data into personal accounts. Move permitted work to an approved tool; never submit customer, staff or supplier personal data to personal accounts.
AI-07	Marketing images — Fictional Image Studio	**Approved with conditions.** The Head of Marketing reviews images before publication.
AI-08	Meeting transcription and summaries — Copilot in Teams	**Approved with conditions.** The meeting organiser checks the transcript and summary before circulation; notice, settings and retention must be confirmed.
AI-09	Sales quotation drafting — Microsoft 365 Copilot	**Approved with conditions.** The Sales Manager or an authorised salesperson checks commercial terms before sending.
AI-10	Supplier invoice data extraction — cloud accounting feature	**Approved with conditions.** Finance checks extracted fields before posting or payment.
AI-11	Customer reorder reminders — marketing automation AI	**Restricted pending consent check.** Do not send reminders until consent and audience rules are checked and approved.
AI-12	Customer web-chat bot — Fictional Helpdesk widget	**Not yet live.** A pilot may proceed only after controls are in place and approved; Customer Service must review pilot responses before any customer-facing use.
AI-13	Outsourced web developer's code assistant	**Approved with conditions.** The developer reviews and tests suggestions before changes are approved. The vendor, settings and repository-access evidence are missing.
AI-14	Demand forecasting — inventory system	**Approved.** The Purchasing Lead reviews forecast exceptions and stock decisions.
AI-15	Supplier-correspondence translation	**Approved with conditions.** A Purchasing user checks material terms before responding or ordering. The tool and its data handling remain to be confirmed.

ID	Tool or use	Status and what you must do
AI-16	B2B credit-limit recommendation — accounting add-on	**Not permitted for automatic decisions.** The Finance Director or authorised finance staff decide; do not allow automatic limit changes.
AI-17	Web-shop fraud flagging — payment provider feature	**Approved with conditions.** Staff review flags before taking adverse action on an order.
AI-18	Delivery route suggestions — Malaysia fulfilment partner system	**Approved.** Partner operations review exceptions and route changes.
AI-19	Internal HR policy Q&A — Microsoft 365 Copilot	**Pilot.** HR checks pilot answers and handles decisions or personal cases.
AI-20	Social media reply drafting — Fictional Draft Studio	**Approved with conditions.** The Head of Marketing or an authorised user checks and posts each reply.

For AI-13 and AI-15, do not infer that a particular tool or data flow is authorised while the vendor, settings or handling remain unknown. Ask the IT Manager or Purchasing Lead, respectively, before entering work information.

4. What you may put into AI tools

The table sets maximum permissions, not an instruction to upload information. **Allowed with conditions** means the specific use, users, data category and settings must be approved and confirmed before entry. Use the least information needed. Remove identifiers where practical and never use a live record merely to test a tool.

An **approved enterprise tool** is a company-approved tool for the relevant use. An **approved vendor feature** is a feature approved for a specific business workflow. **Any other tool** includes personal accounts, trial tools and features not approved for the use.

Information category	Approved enterprise tool	Approved vendor feature	Any other tool
Public information	Allowed, within the approved use	Allowed, within the approved use	Allowed only if no work data is entered and the use does not otherwise breach this policy
Internal information	Allowed with conditions	Allowed with conditions	Not allowed
Confidential information	Allowed with conditions	Allowed with conditions	Not allowed
Customer personal data	Allowed with conditions	Allowed with conditions	Not allowed
Staff personal data	Allowed with conditions	Allowed with conditions	Not allowed
Supplier commercial terms	Allowed with conditions	Allowed with conditions	Not allowed
Payment card data	Not allowed	Not allowed	Not allowed
NRIC numbers	Not allowed	Not allowed	Not allowed

For conditional information, the use owner must confirm that the purpose, access, vendor handling and retention are suitable for the specific workflow. Where the law matters, the Head of Operations should seek confirmation from a Singapore-qualified lawyer.

Do not enter customer lists, CVs, contact details, private meeting content or supplier documents into a tool simply because it appears on an approved list. For example, AI-04 remains on hold, and AI-08's meeting notice and retention controls have not been supplied. Ask before use if the relevant conditions have not been confirmed.

5. Checking outputs before use

AI outputs can be incomplete, outdated, misleading or inappropriate. Before using, sending, publishing or acting on one, make the check required by the register and the workflow.

- **Accuracy:** compare material statements, calculations and summaries with reliable source information. Do not assume a fluent answer is correct.
- **Tone:** check that the content is clear, respectful and suitable for its audience.
- **Personal data:** remove information the recipient does not need, and check that the recipient is authorised to receive what remains.
- **Product claims:** verify product names, specifications, availability and any other claim before publication or communication.
- **Commercial terms:** the Sales Manager or an authorised salesperson must verify quotation terms before sending; Purchasing must check material terms in translated supplier correspondence.
- **Recipients and attachments:** confirm that the right information is going to the right person.
- **Source and completeness:** check whether the output omits a relevant qualification or relies on information you cannot verify.

If you cannot check a material point, do not present it as established fact. Correct or reject the output, or ask the use owner for help. Keep the review or approval record required by your workflow.

6. Prohibited uses

You must not:

- Let AI make a final decision about an individual without human review.
- Use AI-04 to screen or rank candidates while it is on hold.
- Upload customer lists or other work data to an unapproved tool.
- Use a personal AI account for work data, including customer, staff or supplier personal data.
- Impersonate a customer, colleague, supplier or other person.
- Generate or publish fake reviews, testimonials or evidence.
- Submit payment card data or NRIC numbers to an AI tool.
- Allow AI-16 to make an automatic credit-limit decision or change a limit automatically.
- Use AI output to conceal an error, mislead someone or bypass a required approval.
- Enable an unapproved feature, integration or agent permission because it is convenient.
- Use an AI tool to make a decision or take an action outside its approved purpose.

A human review must be real: the reviewer must be able to inspect the relevant information, disagree and stop the action. If the workflow does not allow that, pause it and tell the use owner.

7. Agents and automated actions

No agent may **send, pay, delete or change records** without an approval gate approved under the agent governance checklist. Approval must be given by the person authorised for that action, before it occurs. Do not rely on a suggested control until its operation has been confirmed.

For AI-03, the Purchasing Lead must approve **each purchase order before it is sent**. The proposed agent controls must also set per-order and per-day limits and block orders outside the approved supplier list. The vendor has described the gate and limits but has not yet shown them working in Quillmarrow's own tenant; do not treat them as verified. If they cannot be confirmed, pause sending and escalate to the Purchasing Lead and IT Manager.

Each agent workflow must have a documented way to stop or disable it, commonly called a **kill switch**, and an owner who knows how to use it. The workflow must retain logs sufficient to identify the action, time, approval, outcome and any intervention. The owner and IT Manager should agree how to inspect those records and what to do if logging fails.

If an agent acts unexpectedly, cannot be stopped, exceeds its permission or appears to bypass approval, stop using it where safe to do so and report it immediately under section 10. Do not give it broader access to work around the problem.

8. Customers and transparency

Do not mislead a customer about who or what is handling their request. If a customer-facing web-chat pilot is approved, tell customers clearly that they are interacting with an AI-enabled bot and offer a route to a person. Customer Service must review pilot responses before any customer-facing use, as required for AI-12.

Where AI has materially helped prepare a communication, disclose that assistance when the context, customer request or approved workflow calls for it. Do not claim that an AI output was independently verified unless the relevant check was actually done. Escalate uncertain, sensitive or disputed matters to a person who can respond.

AI-12 is not yet live. Disclosure and hand-off do not by themselves authorise a pilot.

9. New AI tools and features

Before trying a new tool for work, ask the intended use owner and the COO. Describe the purpose, users, information categories, proposed permissions, outputs and decisions or actions involved. A vendor-enabled AI feature counts as a new use even if the underlying business system is already approved.

The use owner should work with the IT Manager and Head of Operations to check relevant settings, data handling, access and vendor evidence. Use the vendor and model due-diligence questionnaire where applicable. Record what was supplied and what remains missing; a vendor statement alone is not proof that Quillmarrow's settings match it.

The use is then entered or updated in the AI use register with its owner, tier, status, conditions and review arrangements. Under Quillmarrow's working method, the use owner approves Tier 1 uses; the use owner and COO approve Tier 2 uses; and the MD approves Tier 3 High uses. A prohibited use must not start until it is redesigned and the appropriate decision is made.

Do not enter work data, connect a system or enable an agent while approval is pending. Approval of a pilot is not approval for wider deployment.

10. Incidents and mistakes

Report an actual or suspected incident **immediately to the Head of Operations, as DPO, and the IT Manager**. Do this even if you corrected the mistake or are unsure whether harm occurred. Do not delete relevant records or attempt to hide the problem.

Examples include entering information into the wrong tool, sending an incorrect AI-drafted reply, link-sharing a folder unexpectedly, exposing a transcript, an agent sending an unapproved order, or discovering that a control or log is not working. Stop the affected activity where safe, preserve the relevant details and follow the IT Manager's containment instructions.

If personal data may be affected, the DPO assesses whether the breach is notifiable. If it is assessed as notifiable, Quillmarrow must notify the PDPC as soon as practicable and no later than **three calendar days after that assessment**. The DPO should confirm legal requirements and the appropriate response with a Singapore-qualified lawyer where needed. Do not wait for a complete investigation before making the internal report.

11. Roles and responsibilities

Role	Responsibility
Managing Director	Approves High-tier uses and provides management oversight.
COO	Owns this policy and the AI programme; coordinates tiering and approval decisions.
Head of Operations and DPO	Coordinates data-protection input and assesses potential personal-data breaches.
Use owners	Define permitted purposes, users, checks and records; review their uses and escalate changes.
IT Manager	Coordinates technical settings, access, incident containment and the managed service provider.
Staff and contractors	Follow approved conditions, check outputs, protect information and report mistakes.
Malaysia fulfilment partner	Follows the agreed arrangements and applicable policy requirements when using Quillmarrow data.
Reviewer	Supports the review and documentation; does not act as Quillmarrow's DPO or approve its operational decisions.

12. Training and acknowledgement

Before using an approved AI workflow, you must receive guidance appropriate to your role and acknowledge that you understand this policy. The use owner should explain the permitted information, required output checks, approval route and how to report a problem.

The COO coordinates policy awareness. Use owners arrange further guidance when a workflow changes or a recurring error shows that staff need clarification. Contractors and the fulfilment partner should receive the relevant requirements through their agreed working arrangements.

13. Breaches of this policy

If you breach this policy, report it immediately rather than trying to resolve it in secret. Quillmarrow will assess the circumstances, contain any risk and decide what corrective action is needed. Depending on the facts and applicable arrangements, this may include additional training, removal of access or action under the relevant employment or contract process. Questions of law should be confirmed with a Singapore-qualified lawyer.

14. Review

The **COO owns this policy** and will arrange a review at least annually. Review it sooner after a material incident or near-miss, a new or significantly changed AI feature, a change in agent permissions, a relevant change in business processes or information handling, or a management decision requiring different controls.

Use owners should also review their own workflows when those triggers occur. Updates must be communicated to affected staff, contractors and partners before they rely on changed instructions.

15. References

- 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).
- Personal Data Protection Act 2012.

Approved by: Managing Director (role) - date to be entered on approval.

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).