


Safe Harbour (Hong Kong) Consultants Limited

The AI Governance Handbook



From the AI inventory to the board summary

2026 edition · free to distribute

What AI governance evidence actually looks like

This handbook answers one concrete question: when a client's procurement team, an external auditor or a board asks about artificial intelligence, which five documents should already exist, what goes into each one, and who owns it.

We prepare these documents for organisations in Hong Kong and Macau. What follows is the working framework we use on engagements. Please take it and use it.

1. Why this is arriving now

In Hong Kong, AI governance is not driven by a single statute. Several lines are moving at once.

- The Privacy Commissioner for Personal Data (PCPD) has issued guidance on advanced technology and generative AI. It is voluntary, but it is cited in complaint handling and investigations.
- The Protection of Critical Infrastructures (Computer Systems) Ordinance (Cap. 653) has been in force since 1 January 2026, and it places statutory duties on designated operators.
- The Hong Kong Monetary Authority and the Securities and Futures Commission have issued circulars on the use of AI. These are supervisory expectations.
- ISO/IEC 42001:2023 is the certifiable AI management system standard, and it is beginning to appear in commercial contracts.
- Where a group has EU operations or users, the EU AI Act applies in phases by risk tier.

For most Hong Kong organisations, however, the first thing to arrive is not a regulator. It is a client procurement questionnaire, asking you to explain within a few working days which AI systems you use, what personal data they touch, and who supervises them.

The point

Regulation moves in years. Procurement questionnaires arrive in days. Prepared documents turn a multi-week scramble into a one-day reply.

2. The five documents

Together these form a minimum viable, audit-defensible AI governance record. They have an order: the inventory comes first, and the other four are built on it.

01 The AI inventory

This is the foundation. Without it, the risk assessment has nothing to assess and the questionnaire cannot be answered.

- System or tool name, vendor and version
- Purpose, and the business process it supports
- A named owner — a role, not just a department
- Data categories processed, flagged for personal data
- Where the data is stored and processed

- Who may use it, and how access is granted
- Go-live date and last review date
- Risk tier

Common gap

The usual hole is that the inventory lists only what IT approved, and misses the services staff subscribed to themselves. The industry calls this shadow AI, and it is often the first thing a questionnaire asks about.

02 Risk and impact assessment

Assess each AI system for the risk it creates — both the impact on individuals and the risk to the organisation. Write both, because regulators and clients look at different things.

- Risk scenarios — what would actually go wrong
- Likelihood and severity
- Which categories of individuals are affected, and roughly how many
- Mitigations already in place
- Residual risk, and who accepts it
- Review date, and the triggers for reassessment

Common gap

Technical risk only — model accuracy, for instance — while the impact on individuals goes unrecorded. That is the gap regulators look for first.

03 Use policy and human review

The policy has to answer two questions: what staff may use, and when a human must decide.

- Acceptable use, and prohibited uses
- An approved tools list, with version and purpose
- Rules on inputs — for example, no client personal data
- The situations where output must be reviewed by a person
- Disclosure requirements, including whether clients are told
- Who reviews, and how the review is recorded

Common gap

A thorough policy that never says who reviews, or at what point. In practice that means nobody reviews.

04 Audit records

The purpose of the record is to be able to produce it within one working day of being asked.

- Approvals for introducing a tool
- Staff acknowledgements of the policy
- Human review records for high-risk uses
- Incident and exception reports, including outcomes
- Records of the periodic inventory and risk reviews

Common gap

Records scattered across email and chat, impossible to assemble inside the deadline. This is where organisations most often lose marks.

05 One-page board summary

The board needs risk and decisions, not configuration detail. One page, no more.

- Which AI is in use, and for what
- The two or three highest risks, and why
- What has already been done
- What remains open, and what needs a board decision
- When it will next be reviewed

Common gap

Handing the technical file to the board. The result is that nobody reads it, and there is no oversight record.

3. What the questionnaire will ask

These are the questions we see most often in procurement questionnaires, due diligence packs and auditor requests, with the document that answers each one.

Typical question	Document	Reply time once prepared
Which AI systems do you use?	01 Inventory	within one working day
Who supervises them?	01, 05	within one working day
What personal data is involved?	01, 02	two to three working days
Where is the data stored?	01	within one working day
How are high-risk uses controlled?	02, 03	two to three working days

Typical question	Document	Reply time once prepared
Is there an incident route?	04	within one working day
How does management oversee this?	05	within one working day

4. A ninety-day plan

If nothing exists yet, this is the order we would work in.

When	What	Output
Weeks 1–2	Inventory every AI system, including services staff subscribed to themselves	Draft AI inventory
Weeks 3–6	Assess risk and impact system by system; define the high-risk uses	Risk and impact assessment
Weeks 7–10	Publish the use policy and approved tools list; name the reviewers	Use policy and review rules
Weeks 11–13	Stand up the records and the incident route; brief the board	Audit records and board summary
After that	Review the inventory and risks quarterly; review everything annually	Review records

5. Third parties and vendors

Most AI risk does not come from a model you built. It comes from a service you use, and questionnaires tend to probe exactly that.

- Whether the vendor explains how the model is trained and how data is handled
- Whether your data is used for training, and whether that can be switched off
- The jurisdictions where data is stored and processed
- Service levels, breach notification obligations and exit terms
- The vendor's own certifications, such as ISO/IEC 27001 or ISO/IEC 42001
- A list of sub-processors

In practice

Ask these questions during procurement and you get complete answers. Ask afterwards and you usually get a partial reply.

6. Incident handling in five steps

AI incidents do not look like classic security incidents. Biased output, data leakage and a bypassed safeguard all qualify. This is the sequence we recommend.

- 1. Contain — suspend the system or the use case, and preserve the state.
- 2. Scope — which personal data, which clients, how many people.
- 3. Notify — establish whether the PCPD, HKMA, SFC or a client contract imposes a notification duty, and by when.
- 4. Remedy and review — fix the cause, reassess the risk, update the inventory.
- 5. Record — write the four steps above into a single auditable record. That record is itself part of document 04.

Note

Notification clocks usually start at discovery, not at confirmation of cause. Notify first and supplement afterwards.

7. Macau

Macau has no separate AI statute. AI and privacy material published by the PDPB is expressly marked non-binding, but it indicates where supervisory expectations sit. Law 8/2005 (the Personal Data Protection Act) applies to personal data processing in Macau.

Where you operate or have clients in Macau, handle it with Hong Kong: the same five documents, extended with Macau's data processing and notification requirements.

8. Common pitfalls

- Starting with the policy rather than the inventory. The policy ends up vague because it has no subject.
- Treating AI governance as an IT matter. The questionnaire asks about the organisation, and whoever answers needs management authority.
- Aiming for perfection. A one-page version published this month beats a thorough one published next year.
- Covering only generative AI, and missing the predictive models and automated decisions already in production.
- A policy with no review records, which cannot be shown to have been applied.

9. How this relates to ISO/IEC 42001

ISO/IEC 42001:2023 is the certifiable AI management system standard. The five documents above cover most of the operational core of the standard. Certification additionally requires the management system elements: organisational context, leadership commitment, internal audit and management review.

If the goal is answering client questionnaires, the five documents are usually enough. Whether to certify is a separate decision, and it does not have to be tied to questionnaire readiness.

Where we draw the line

We are a readiness and advisory partner. We do not perform certification audits and we do not issue certificates. Certification audits are conducted by an accredited certification body.

10. Glossary

Term	Meaning
AI system	A system that uses machine learning or other automated methods to produce output — predictions, recommendations or content — used in the business.
Shadow AI	AI tools and services used or subscribed to by staff without organisational approval.
Risk tier	A banding of systems by severity and likelihood, used to decide how much control is required.
Human review	A person with authority examines and decides before automated output affects an individual or the organisation.
Management system	Policies, procedures, responsibilities and records organised so that a requirement is met consistently rather than from memory.
Certification body	An accredited body that conducts third-party audits of management systems and issues certificates.

11. Choosing between 42001, 27001 and 27701

The three are often conflated. Their scope is actually clean.

- ISO/IEC 27001 is the information security management system. It addresses security risk, and it is usually the starting point.
- ISO/IEC 27701 extends that to privacy information management — the duties of data controllers and processors.
- ISO/IEC 42001 is the AI management system. It addresses the risks specific to AI: data provenance, models, impact assessment and human oversight.

Where a client asks for one, most organisations start with 27001. Where a client asks specifically for AI governance, 42001 leads, and running the two together shares a great deal of documentation — typically 20 to 30 per cent extra effort rather than double.

Appendix A: an inventory template

These are the fields to carry in your inventory. The first two rows are worked examples, not part of your actual record.

System / vendor	Purpose and owner	Data and location
Example: Microsoft 365 Copilot	Drafting internal documents · Head of IT	Internal documents, no personal data · HK and Singapore
Example: customer service chatbot	Answering customer enquiries · Head of Customer Service	Name, contact details, conversation logs · Hong Kong

Also record the go-live date, the last review date, the risk tier, and whether the tool appears on the approved tools list.

Appendix B: a model answer

The following can be adapted directly when a questionnaire asks how your organisation uses AI.

Template

The organisation maintains an inventory of AI systems recording the tools in use, their purpose, a named owner and the data categories processed. All tools are approved through an approved tools list and governed by a written use policy. Outputs involving personal data, or which affect individuals, require review by a designated person before use. The inventory and risk assessment are reviewed quarterly, and the reviews are recorded. The inventory and a summary of the risk assessment can be provided within one working day.

Appendix C: use policy clauses

These clauses can be adopted or adapted directly. Publish them alongside the HR handbook, and ask staff to acknowledge them.

- 1. Scope. This policy applies to all employees, contractors and agents, and to every AI tool they use in the course of business.
- 2. Approved tools. Staff may use only the tools on the approved tools list. The list is reviewed quarterly by IT with the data protection lead.
- 3. Prohibited uses. No use to make automated decisions with a legal or similarly significant effect on an individual. No use to generate or circulate discriminatory, misleading or infringing content.
- 4. Inputs. Client or employee personal data, confidential material, unpublished financial information and third-party intellectual property must not be entered without written approval and a completed risk assessment.
- 5. Human review. Where output will affect a client, an employee or a third party, or is intended for publication, a designated person must review it before use and record the review.
- 6. Disclosure. On request, the organisation will describe the AI systems it uses and how data is handled.

- 7. Incident reporting. Staff must notify IT and the data protection lead within the same working day of any anomalous output, bias or data exposure.
- 8. Breach. A breach of this policy may be a disciplinary matter.

Note

Follow the clauses with one page of who does what, naming the specific action for each role. Without it the clauses cannot be applied.

Appendix D: a board summary template

Item	Content
Use	X AI systems in use this quarter, up or down Y on last quarter.
Top risks	Two or three, with a short reason for each.
Actions taken	Controls completed, and what they actually achieved.
Open items	What needs a board decision, with a recommended option.
Compliance	Position against applicable requirements, and readiness for ISO/IEC 42001.
Next review	Date and owner.

One page, no more. If it needs a second page, the decision you want from the board is probably not yet clear.

12. About Safe Harbour

Safe Harbour (Hong Kong) Consultants Limited is a Hong Kong consultancy working on AI governance, ISO certification and IT audit, for organisations in Hong Kong and Macau.

We publish our consultants' credentials, our method, and the source and date behind every regulatory position we take.

- Telephone: +852 5382 0328
- Email: info@safeharbour.hk
- WhatsApp: +852 5382 0328
- Address: Unit A, 22/F, 235 Commercial Centre, 235 Wing Lok Street, Sheung Wan, Hong Kong

This handbook is general information, not legal advice. Regulatory positions change; verify the current version before relying on it.