CYBERGLASS AI GOVERNANCE KIT

AI Strategy and Governance Charter

Sets [Company Name]'s direction for using artificial intelligence, the principles we follow, how AI risk is tiered and who decides.

Organization	[Company Name]
Document ID	AI-01
Document owner	[Name, title]
Approved by	[Name, title]
Version	1.0 [draft / approved]
Effective date	[YYYY-MM-DD]
Next review	[YYYY-MM-DD] (at least every 12 months)
Classification	Internal

PREVIEW

About this template (delete this box before publishing)

This document is a template to help [Company Name] write its own policy. It is a guideline based on common industry practice and publicly available frameworks, not a finished policy, legal advice or a guarantee of compliance. Every organization is different: review each requirement, remove what does not apply, add what is missing, and set values your organization can actually meet. Have the final version reviewed by people who know your systems, obligations and risks.

- Write this with the business leaders who want to use AI, not only with IT and legal. Strategy without use cases goes nowhere.
- Keep the risk tiers (section 6) simple. Most organizations need only four.
- Approve it at executive level, then use AI-02 to AI-06 and the workbook to put it into practice.
- Replace every highlighted **[placeholder]** with your organization's details, then remove the highlighting.
- Grey **Guidance** boxes explain what auditors, insurers and clients look for. Read them, then delete them before publishing.

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

Executive summary

IN PLAIN WORDS *for leaders, clients and procurement teams (keep and adapt)*

What it does: Sets why and how [Company Name] uses AI, the principles we follow, four risk tiers and who decides.

How it protects clients: No AI system that affects clients or their data goes live without an owner, a risk tier and approval at the right level.

What clients can ask for: Our AI principles, the AI systems used on their work, and the client assurance statement in Appendix A.

Measured by: AI systems with owner and tier (target 100%), assessments current, training completion and AI incidents (section 10).

Real-world context: why this document matters

When	What happened	What it means for this document
Jul 2025	The global average cost of a breach was US$4.44M. 13% of organizations reported breaches of AI models or applications, and 97% of those said they lacked proper AI access controls.	Most organizations breached through AI lacked AI access controls. Lifecycle controls (section 9) and decision rights (section 8) close that gap.
Feb 2024	A Canadian tribunal held an airline liable for its website chatbot's wrong advice about bereavement fares. It rejected the argument that the chatbot was responsible for its own statements.	The organization is accountable for what its AI says. Principle 'human accountability' (section 4) is a legal reality.
Jan 2024	A finance employee in the Hong Kong office of a global engineering firm transferred about US$25M (HK$200M) after a video call in which the 'CFO' and other colleagues were all deepfakes.	AI also changes the threat. The strategy should cover AI-enabled fraud, not only your own AI use.

Summarized by CyberGlass from public reporting: IBM Cost of a Data Breach Report 2025; BC Civil Resolution Tribunal decision, 14 Feb 2024; CNN and SCMP, May 2024. Organizations are described, not named. Figures are as publicly reported at the time.

PREVIEW

How this document connects

Document or register	How records and decisions flow
AI-02 Acceptable Use Policy	Staff rules under this charter
AI-03 Risk Method	Applies the risk tiers in section 6
AI-04 Tool Approval	Implements decision rights in section 8
IT-03 Governance Framework	AI Governance Committee reports to the Security Steering Committee or board
AI Governance Workbook	AI Inventory, AI Risk Register (AR-NNN), Tool Requests (TR-NNN), Shadow AI Findings (SF-NNN), Control Register
IT-06 Incident Response Plan	All AI incidents use the same INC-YYYY-NNN log and severity scale
Kit risk register (IT/OT/Rail workbook)	AI risks rated High or Critical are copied as enterprise risks R-NNN

Record flow example: AI-06 review finds 46 staff using a free chatbot with customer data (SF-001) → reported as INC-2026-031 under AI-05 playbook 5.6 and IT-06 → AI risk AR-009 'customer data leakage via unapproved AI' (L4 × I4 = High) → enterprise risk R-018 updated → decision 'Replace' with the approved assistant (TR-006) → KPI 'open shadow AI findings' tracked monthly.

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

CYBERGLASS AI GOVERNANCE KIT

AI Acceptable Use Policy

Rules for how everyone at [Company Name] may use AI tools, assistants and agents, and what data may go into them.

Organization	[Company Name]
Document ID	AI-02
Document owner	[Name, title]
Approved by	[Name, title]
Version	1.0 [draft / approved]
Effective date	[YYYY-MM-DD]
Next review	[YYYY-MM-DD] (at least every 12 months)
Classification	Internal

PREVIEW

About this template (delete this box before publishing)

This document is a template to help [Company Name] write its own policy. It is a guideline based on common industry practice and publicly available frameworks, not a finished policy, legal advice or a guarantee of compliance. Every organization is different: review each requirement, remove what does not apply, add what is missing, and set values your organization can actually meet. Have the final version reviewed by people who know your systems, obligations and risks.

- Publish the approved tools list (section 4) at the same time as the policy. Banning tools without offering alternatives drives shadow AI.
- Adapt the data table (section 5) to your own classification scheme in IT-04.
- Collect acknowledgements from all staff and contractors, and repeat yearly.
- Replace every highlighted **[placeholder]** with your organization's details, then remove the highlighting.
- Grey **Guidance** boxes explain what auditors, insurers and clients look for. Read them, then delete them before publishing.

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

Executive summary

IN PLAIN WORDS *for leaders, clients and procurement teams (keep and adapt)*

What it does: Tells every member of staff which AI tools they may use and what information may go into them.

How it protects clients: Client and personal data only goes into approved enterprise AI tools that do not train on it; secrets never go into any AI tool.

What clients can ask for: The approved AI tools list and the data rules in section 5.

Measured by: Staff acknowledgement and AI training completion; AI misuse reported and handled.

Real-world context: why this document matters

When	What happened	What it means for this document
Apr 2023	Engineers at a large electronics manufacturer pasted confidential source code and meeting notes into a public AI chatbot. The company then restricted generative AI tools on company devices.	Section 5 data rules: confidential data never goes into unapproved tools.
Jan 2024	A finance employee in the Hong Kong office of a global engineering firm transferred about US$25M (HK$200M) after a video call in which the 'CFO' and other colleagues were all deepfakes.	Section 13: staff must report suspected AI impersonation immediately.
Jul 2025	The global average cost of a breach was US$4.44M. 13% of organizations reported breaches of AI models or applications, and 97% of those said they lacked proper AI access controls.	Unapproved AI use was linked with higher breach costs; give staff approved tools (section 4).

Summarized by CyberGlass from public reporting: Public reporting, May 2023; CNN and SCMP, May 2024; IBM Cost of a Data Breach Report 2025. Organizations are described, not named. Figures are as publicly reported at the time.

How this document connects

Document or register	How records and decisions flow
AI-01 Charter	Principles and tiers
AI-04 Tool Approval	How tools reach the approved list
AI-05 Incident Playbook	What happens after a report under section 13
IT-04 Policy 7.6	Parent policy rule
AI Governance Workbook	AI Inventory, AI Risk Register (AR-NNN), Tool Requests (TR-NNN), Shadow AI Findings (SF-NNN), Control Register
IT-06 Incident Response Plan	All AI incidents use the same INC-YYYY-NNN log and severity scale
Kit risk register (IT/OT/Rail workbook)	AI risks rated High or Critical are copied as enterprise risks R-NNN

Record flow example: AI-06 review finds 46 staff using a free chatbot with customer data (SF-001) → reported as INC-2026-031 under AI-05 playbook 5.6 and IT-06 → AI risk AR-009 'customer data leakage via unapproved AI' (L4 × I4 = High) → enterprise risk R-018 updated → decision 'Replace' with the approved assistant (TR-006) → KPI 'open shadow AI findings' tracked monthly.

PREVIEW

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

CYBERGLASS AI GOVERNANCE KIT

AI Risk and Impact Assessment Method

How [Company Name] tiers AI systems, assesses their risks and their impact on people, tests them and decides whether they can be used.

Organization	[Company Name]
Document ID	AI-03
Document owner	[Name, title]
Approved by	[Name, title]
Version	1.0 [draft / approved]
Effective date	[YYYY-MM-DD]
Next review	[YYYY-MM-DD] (at least every 12 months)
Classification	Internal

PREVIEW

About this template (delete this box before publishing)

This document is a template to help [Company Name] write its own policy. It is a guideline based on common industry practice and publicly available frameworks, not a finished policy, legal advice or a guarantee of compliance. Every organization is different: review each requirement, remove what does not apply, add what is missing, and set values your organization can actually meet. Have the final version reviewed by people who know your systems, obligations and risks.

- Use the short tiering questions (section 4) for every request; do the full assessment only for Medium and High tiers.
- Record results in the AI workbook (Inventory and AI Risk Register sheets) so they can be reported.
- Involve privacy and the people who will use the output. They spot problems technical teams miss.
- Replace every highlighted **[placeholder]** with your organization's details, then remove the highlighting.
- Grey **Guidance** boxes explain what auditors, insurers and clients look for. Read them, then delete them before publishing.

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

Executive summary

IN PLAIN WORDS *for leaders, clients and procurement teams (keep and adapt)*

What it does: A repeatable way to rate AI risks, including impact on people, and test AI before use.

How it protects clients: AI that affects clients or processes their data is tested for accuracy, bias, data leakage and prompt injection before go-live.

What clients can ask for: A summary of the assessment and test results for AI used on their work.

Measured by: Medium and High-tier systems with a current assessment (target 100%).

Real-world context: why this document matters

When	What happened	What it means for this document
Feb 2024	A Canadian tribunal held an airline liable for its website chatbot's wrong advice about bereavement fares. It rejected the argument that the chatbot was responsible for its own statements.	Accuracy risk has legal consequences; rate it with legal input (section 5).
Jul 2025	The global average cost of a breach was US$4.44M. 13% of organizations reported breaches of AI models or applications, and 97% of those said they lacked proper AI access controls.	Security of AI systems needs access controls; test them (section 8).

Summarized by CyberGlass from public reporting: BC Civil Resolution Tribunal decision, 14 Feb 2024; IBM Cost of a Data Breach Report 2025. Organizations are described, not named. Figures are as publicly reported at the time.

How this document connects

Document or register	How records and decisions flow
IT-05 Risk Method	Same 5×5 scales so AI risks compare with other risks
AI-04 Vendor Assessment	Vendor score feeds supply chain risk
AI-01 Charter section 8	Approval authority by tier
AI Governance Workbook	AI Inventory, AI Risk Register (AR-NNN), Tool Requests (TR-NNN), Shadow AI Findings (SF-NNN), Control Register
IT-06 Incident Response Plan	All AI incidents use the same INC-YYYY-NNN log and severity scale
Kit risk register (IT/OT/Rail workbook)	AI risks rated High or Critical are copied as enterprise risks R-NNN

Record flow example: AI-06 review finds 46 staff using a free chatbot with customer data (SF-001) → reported as INC-2026-031 under AI-05 playbook 5.6 and IT-06 → AI risk AR-009 'customer data leakage via unapproved AI' (L4 × I4 = High) → enterprise risk R-018 updated → decision 'Replace' with the approved assistant (TR-006) → KPI 'open shadow AI findings' tracked monthly.

PREVIEW

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

CYBERGLASS AI GOVERNANCE KIT

AI Tool Approval and Vendor Assessment

The intake process, request form, vendor security and privacy questionnaire and scoring used to approve AI tools and AI vendors at [Company Name].

Organization	[Company Name]
Document ID	AI-04
Document owner	[Name, title]
Approved by	[Name, title]
Version	1.0 [draft / approved]
Effective date	[YYYY-MM-DD]
Next review	[YYYY-MM-DD] (at least every 12 months)
Classification	Internal

PREVIEW

About this template (delete this box before publishing)

This document is a template to help [Company Name] write its own policy. It is a guideline based on common industry practice and publicly available frameworks, not a finished policy, legal advice or a guarantee of compliance. Every organization is different: review each requirement, remove what does not apply, add what is missing, and set values your organization can actually meet. Have the final version reviewed by people who know your systems, obligations and risks.

- Make the request form easy. If approval takes months, people will use unapproved tools instead.
- Use the questionnaire for vendors whose AI will process Internal or higher data; Low-tier public-data tools may need only the request form.
- Contract points in section 6 are a checklist for your legal team, not contract wording.
- Replace every highlighted **[placeholder]** with your organization's details, then remove the highlighting.
- Grey **Guidance** boxes explain what auditors, insurers and clients look for. Read them, then delete them before publishing.

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

Executive summary

IN PLAIN WORDS *for leaders, clients and procurement teams (keep and adapt)*

What it does: How AI tools and vendors are requested, assessed and approved, including a scored vendor questionnaire.

How it protects clients: AI vendors are checked for data use, retention, security, location and incident notice before client data is shared.

What clients can ask for: The list of AI vendors that process their data and the shared responsibility table in Appendix A.

Measured by: Vendors assessed before use (target 100%); request turnaround time.

Real-world context: why this document matters

When	What happened	What it means for this document
May 2023	A criminal group mass-exploited a zero-day flaw in a widely used managed file-transfer product, stealing data from thousands of organizations. Many victims learned of it from their own suppliers.	Supplier software can be the breach; AI vendors need the same scrutiny.
Jul 2025	The global average cost of a breach was US$4.44M. 13% of organizations reported breaches of AI models or applications, and 97% of those said they lacked proper AI access controls.	Access controls on AI services are the gap to check (questionnaire 'Access and identity').

Summarized by CyberGlass from public reporting: Public reporting and government advisories, 2023; IBM Cost of a Data Breach Report 2025. Organizations are described, not named. Figures are as publicly reported at the time.

How this document connects

Document or register	How records and decisions flow
AI Workbook Vendor Assessment and Tool Requests	Scores and decisions
AI-02 approved list	Updated after decisions
Workbook Suppliers sheet	AI vendors added with review dates
AI Governance Workbook	AI Inventory, AI Risk Register (AR-NNN), Tool Requests (TR-NNN), Shadow AI Findings (SF-NNN), Control Register
IT-06 Incident Response Plan	All AI incidents use the same INC-YYYY-NNN log and severity scale
Kit risk register (IT/OT/Rail workbook)	AI risks rated High or Critical are copied as enterprise risks R-NNN

Record flow example: TR-003 AI note-taker → vendor score 55% with 0 in Data use → declined (section 5) → approved alternative offered → request closed in 7 days.

PREVIEW

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

CYBERGLASS AI GOVERNANCE KIT

AI Incident Response Playbook

How [Company Name] recognizes and responds to AI-related incidents, from data entered into the wrong tool to manipulated agents and deepfake impersonation.

Organization	[Company Name]
Document ID	AI-05
Document owner	[Name, title]
Approved by	[Name, title]
Version	1.0 [draft / approved]
Effective date	[YYYY-MM-DD]
Next review	[YYYY-MM-DD] (at least every 12 months)
Classification	Internal

PREVIEW

About this template (delete this box before publishing)

This document is a template to help [Company Name] write its own policy. It is a guideline based on common industry practice and publicly available frameworks, not a finished policy, legal advice or a guarantee of compliance. Every organization is different: review each requirement, remove what does not apply, add what is missing, and set values your organization can actually meet. Have the final version reviewed by people who know your systems, obligations and risks.

- Use with the Incident Response Plan (IT-06). IT-06 sets the overall process; this playbook covers what is different for AI.
- Make sure staff know that reporting a mistake quickly is welcome; most AI incidents start as an honest error.
- Exercise the deepfake impersonation playbook with finance and executive assistants.
- Replace every highlighted **[placeholder]** with your organization's details, then remove the highlighting.
- Grey **Guidance** boxes explain what auditors, insurers and clients look for. Read them, then delete them before publishing.

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

Executive summary

IN PLAIN WORDS *for leaders, clients and procurement teams (keep and adapt)*

What it does: How AI incidents are handled: data in the wrong tool, prompt injection, goal hijack and memory poisoning, runaway agents, unregistered agents, harmful output, vendor breaches and deepfake fraud.

How it protects clients: Clear service levels for triage, containment and client notification (Appendix A).

What clients can ask for: Our notification commitment and how they will be kept informed.

Measured by: Time to triage, contain and notify, against the targets in Appendix A.

Real-world context: why this document matters

When	What happened	What it means for this document
Jan 2024	A finance employee in the Hong Kong office of a global engineering firm transferred about US$25M (HK$200M) after a video call in which the 'CFO' and other colleagues were all deepfakes.	Playbook 5.5 exists because a video call fooled a trained finance employee.
Apr 2023	Engineers at a large electronics manufacturer pasted confidential source code and meeting notes into a public AI chatbot. The company then restricted generative AI tools on company devices.	Playbook 5.1: data entered into the wrong tool.
Feb 2024	A Canadian tribunal held an airline liable for its website chatbot's wrong advice about bereavement fares. It rejected the argument that the chatbot was responsible for its own statements.	Playbook 5.3: wrong output reached customers.

Summarized by CyberGlass from public reporting: CNN and SCMP, May 2024; Public reporting, May 2023; BC Civil Resolution Tribunal decision, 14 Feb 2024. Organizations are described, not named. Figures are as publicly reported at the time.

How this document connects

Document or register	How records and decisions flow
AI-06 Shadow AI Review	Findings can trigger playbook 5.6
AI-03 Risk Method	Reassessment after incidents (section 8)
AI Governance Workbook	AI Inventory, AI Risk Register (AR-NNN), Tool Requests (TR-NNN), Shadow AI Findings (SF-NNN), Control Register
IT-06 Incident Response Plan	All AI incidents use the same INC-YYYY-NNN log and severity scale
Kit risk register (IT/OT/Rail workbook)	AI risks rated High or Critical are copied as enterprise risks R-NNN

Record flow example: AI-06 review finds 46 staff using a free chatbot with customer data (SF-001) → reported as INC-2026-031 under AI-05 playbook 5.6 and IT-06 → AI risk AR-009 'customer data leakage via unapproved AI' (L4 × I4 = High) → enterprise risk R-018 updated → decision 'Replace' with the approved assistant (TR-006) → KPI 'open shadow AI findings' tracked monthly.

PREVIEW

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

CYBERGLASS AI GOVERNANCE KIT

Shadow AI Baseline Review Guide

A practical, proportionate method to find out which AI tools people at [Company Name] actually use, with which data, and what to do about it.

Organization	[Company Name]
Document ID	AI-06
Document owner	[Name, title]
Approved by	[Name, title]
Version	1.0 [draft / approved]
Effective date	[YYYY-MM-DD]
Next review	[YYYY-MM-DD] (at least every 12 months)
Classification	Internal

PREVIEW

About this template (delete this box before publishing)

This document is a template to help [Company Name] write its own policy. It is a guideline based on common industry practice and publicly available frameworks, not a finished policy, legal advice or a guarantee of compliance. Every organization is different: review each requirement, remove what does not apply, add what is missing, and set values your organization can actually meet. Have the final version reviewed by people who know your systems, obligations and risks.

- Start with one team or department and one month of data. A small, accurate baseline beats a large, vague one.
- Agree the method with HR, Legal and Privacy before collecting any data about staff activity.
- Treat findings as information, not as grounds for blame. The goal is to give people approved tools that meet their needs.
- Replace every highlighted **[placeholder]** with your organization's details, then remove the highlighting.
- Grey **Guidance** boxes explain what auditors, insurers and clients look for. Read them, then delete them before publishing.

Preview pages · full editable Word/Excel files at cyberglass.ai/shop

Executive summary

IN PLAIN WORDS *for leaders, clients and procurement teams (keep and adapt)*

What it does: Finds which AI tools staff really use and which AI agents or automations teams have built without approval, with which data, and decides what to approve, replace or block.

How it protects clients: Unapproved tools that could expose client data are found and dealt with, at least every six months.

What clients can ask for: Confirmation that a shadow AI review has been done and its main actions.

Measured by: Open shadow AI findings (target: falling); controls tested with synthetic data.

Real-world context: why this document matters

When	What happened	What it means for this document
Apr 2023	Engineers at a large electronics manufacturer pasted confidential source code and meeting notes into a public AI chatbot. The company then restricted generative AI tools on company devices.	Why discovery matters: staff use what is easiest.
Jul 2025	The global average cost of a breach was US$4.44M. 13% of organizations reported breaches of AI models or applications, and 97% of those said they lacked proper AI access controls.	Shadow AI and missing access controls were linked with higher breach costs.

Summarized by CyberGlass from public reporting: Public reporting, May 2023; IBM Cost of a Data Breach Report 2025. Organizations are described, not named. Figures are as publicly reported at the time.

How this document connects

Document or register	How records and decisions flow
AI-05 Incident Playbook 5.6	Widespread shadow AI with sensitive data is handled as an incident
AI-04 Tool Approval	Decisions 'Approve' or 'Replace' create tool requests
AI-02 Policy	Training and approved list updates
AI Governance Workbook	AI Inventory, AI Risk Register (AR-NNN), Tool Requests (TR-NNN), Shadow AI Findings (SF-NNN), Control Register
IT-06 Incident Response Plan	All AI incidents use the same INC-YYYY-NNN log and severity scale
Kit risk register (IT/OT/Rail workbook)	AI risks rated High or Critical are copied as enterprise risks R-NNN

Record flow example: AI-06 review finds 46 staff using a free chatbot with customer data (SF-001) → reported as INC-2026-031 under AI-05 playbook 5.6 and IT-06 → AI risk AR-009 'customer data leakage via unapproved AI' (L4 × I4 = High) → enterprise risk R-018 updated → decision 'Replace' with the approved assistant (TR-006) → KPI 'open shadow AI findings' tracked monthly.

PREVIEW

Preview pages · full editable Word/Excel files at cyberglass.ai/shop