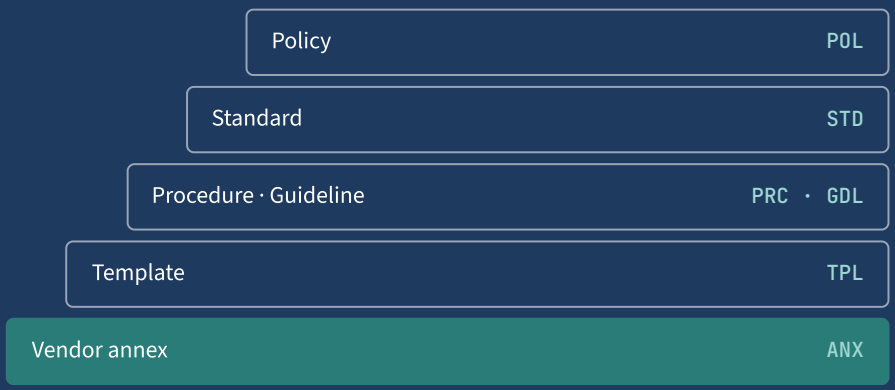
Anchor Sprint
AI DIGITAL AGENCY

PRE-READ

AI GOVERNANCE KIT (AGK) · MALAYSIA & SINGAPORE

Govern how your people use AI, and how your teams build it

A ready-to-adopt set of policies, standards, procedures and templates. It works with any AI vendor, scales from a 50-person SME to a regulated bank, and maps to Malaysian, Singaporean and international frameworks.



48
templates in Word and Excel

422
auditable controls, each with an ID

3
adoption tiers: Lite, Standard, Regulated

4
vendor annexes: Claude, OpenAI, Microsoft, Google

At a glance

What it is. Templates for the rules, roles and records an organisation needs to use and build AI safely.

Who it's for. Organisations from about 50 staff to regulated banks in Malaysia and Singapore, on any AI platform.

What to do first. Take the two-minute check below, then open the first document it names.

Contents

PART 1 **Get started**

START HERE
Which path are you on? 2

TWO-MINUTE CHECK
Find your tier and your first five documents 2

READING ORDER
Read these in this order 3

FOR LEADERSHIP
Six decisions to make in week 1 3

PART 2 **Adopt in 90 days**

HOW TO ADOPT
The first 90 days for an organisation starting out 4

TERMS YOU'LL SEE
Ten terms used across the kit 5

DOCUMENT MAP
What each tier adopts: all 48 documents 5

PART 3 **How the kit is built**

HOW IT'S BUILT
Five levels, from board principle to platform setting 7

FRAMEWORKS BEHIND THE KIT
Built on ISO/IEC 42001, risk-managed with NIST, mapped to local rules 8

TWO TRACKS, ONE CORE
Using AI and building AI are governed separately 9

SHARED RULES
One vocabulary across every document 9

DAY-TO-DAY FLOWS
What the kit looks like in practice 10

WHO DOES WHAT
Named people, not committees alone 11

PART 4 **Context and support**

WHY NOW · STATUS AS AT OCTOBER 2026
AI rules in Malaysia and Singapore are moving from voluntary to expected 12

QUESTIONS LEADERS ASK
FAQ 13

WORKING WITH THE TEMPLATES
What to keep, what to change 13

NEXT STEP
From free templates to a governed AI rollout 14

PART 1 **Get started**

START HERE

Which path are you on?

Pick the description closest to you, or take the two-minute check below. Each path tells you which document to open first.

Starting out with AI

Staff use, or want to use, vendor AI tools. Little or no AI policy yet.

Open AGK-001 Lite Quick Start → follow the 90-day plan

Using AI tools that act

Connectors to email or files, browser agents, or shared skills are on, or about to be.

AGK-001, then Track A: AGK-STD-A03 and A04 before switching features on

Building agents or customer-facing AI

You build AI applications or agents, or AI touches customers or decisions about people.

AGK-000, then Track B (AGK-STD-B01 to B06) and AGK-PRC-02

Regulated or seeking ISO/IEC 42001

Financial institution under BNM or MAS, healthcare, EU exposure, or certification planned.

AGK-000, then the Statement of Applicability in AGK-TPL-08

TWO-MINUTE CHECK

Find your tier and your first five documents

#	Question	If yes
1	Are you regulated by BNM, MAS or a health regulator, have EU exposure, or plan ISO/IEC 42001 certification?	**Regulated tier**
2	Do you build AI agents or AI applications, in-house or with a partner?	**Standard tier (unless question 1 makes you Regulated)**
3	Do customers interact with your AI, or does AI help make decisions about people (hiring, credit, pricing)?	**Standard tier (unless question 1 makes you Regulated)**
4	Do more than 250 staff use AI tools?	**Standard tier (unless question 1 makes you Regulated)**
5	Are connectors that can write, browser agents or shared skills switched on, or planned?	**Adopt AGK-STD-A03 and A04 before switching these features on**
6	Do you already have an approved AI tool and an AI usage policy?	**Compare your policy and tool with AGK-POL-02 and AGK-STD-A01 rather than starting again**

Lite

No to questions 1–4: vendor AI tools, no in-house builds.

1. `AGK-001` Lite Quick Start
2. `AGK-PRC-07` AI Discovery & Shadow AI
3. `AGK-POL-02` Acceptable Use Policy (with AGK-GDL-A01 and AGK-TPL-07)
4. `AGK-STD-A02` AI Data Handling
5. `AGK-STD-A01` Platform Baseline + your Vendor Annex

Standard

Yes to any of questions 2–4.

1. `AGK-000` Framework & Document Architecture
2. `AGK-PRC-07` AI Discovery & Shadow AI
3. `AGK-POL-01` Governance Charter
4. `AGK-POL-02` Acceptable Use Policy
5. `AGK-STD-A01` Platform Baseline + your Vendor Annex

Regulated

Yes to question 1.

1. `AGK-000` Framework & Document Architecture
2. `AGK-TPL-08` Control Library and ISO/IEC 42001 SoA (set tier to Regulated)
3. `AGK-POL-01` Governance Charter
4. `AGK-PRC-02` Risk & Impact Assessment
5. `AGK-PRC-03` Vendor & Model Assessment

READING ORDER

Read these in this order

The order matters: each document assumes the ones before it. Times are for a first read.

1	This pre-read	Everyone involved	10 min
2	AGK-001 Lite Quick Start (Standard and Regulated: AGK-000 Framework)	AI sponsor, AI Governance Lead	15 min
3	AGK-PRC-07 AI Discovery & Shadow AI	AI Governance Lead, IT	10 min
4	AGK-POL-02 Acceptable Use Policy and AGK-GDL-A01 Safe AI Use Guide	HR, AI Governance Lead	20 min
5	AGK-STD-A02 AI Data Handling (the data × tool matrix)	DPO, IT Security	10 min
6	AGK-STD-A01 Platform Baseline and the Vendor Annex for your platform	AI Platform Administrator	30 min
7	AGK-GDL-C02 Adoption Roadmap and AGK-GDL-C03 Change & Communication	AI sponsor, HR	25 min

FOR LEADERSHIP

Six decisions to make in week 1

Nothing else moves until these are made. Each takes one conversation, not a project.

Who owns AI?

Name an AI sponsor (executive) and an AI Governance Lead.

AGK-POL-01

Who decides?

Form a three-person committee: sponsor, Governance Lead, DPO.

AGK-POL-01

How much risk is acceptable?

One paragraph of risk appetite: what AI may and may never be used for.

AGK-POL-01

Which AI tool is approved?

Confirm the platform(s) staff should use instead of personal accounts.

AGK-GDL-C01

Time and budget

About half a day a week for the Governance Lead in the first 90 days, plus licences.

AGK-001

What we say about jobs

A short, honest statement on how AI affects roles, before staff ask.

AGK-GDL-C03

PART 2 **Adopt in 90 days**

HOW TO ADOPT

The first 90 days for an organisation starting out

This is the Lite path from AGK-001. Each step ends with something approved, configured or populated. Standard and Regulated adopters follow the same order and add the remaining documents.

Week 1 — **Name the owners**

AI sponsor, AI Governance Lead and a three-person committee. Agree the risk appetite in one paragraph.

`AGK-001 · POL-01`

Weeks 1–2 — **Find the AI already in use**

Staff survey, expense data, network and SaaS discovery, with a 30-day amnesty for self-declared tools.

`AGK-PRC-07 · TPL-01`

Week 2 — **Set the rules**

Approve the Acceptable Use Policy and data rules. Every user acknowledges before access.

`AGK-POL-02 · STD-A02 · GDL-A01 · TPL-07`

Weeks 2–3 — **Configure the platform**

Apply the baseline with the vendor annex. Block personal AI accounts on company devices.

`AGK-STD-A01 · ANX-*`

Weeks 3–4 — **Communicate and train**

Key messages, staff FAQ, champions, a safe channel for concerns, fundamentals training.

`AGK-GDL-C03 · STD-C01`

Weeks 4–6 — **Pick the first use cases**

Run a two-hour workshop, score ideas on value, feasibility and risk, pilot two or three quick wins.

`AGK-GDL-C02 · TPL-12`

Weeks 4–6 — **Be ready for problems**

Intake form, incident log, who to call, and the 72-hour breach clock.

`AGK-PRC-01 · PRC-04 · TPL-06`

Weeks 6–12 — **Pilot, measure, decide**

Track benefits. Adopt connector or browser standards before switching those features on.

`AGK-TPL-12 · STD-A03 / A04`

Day 90 — **Review**

Committee reviews what is in use, incidents, benefits and next quarter. Move to Standard when you build agents or scale.

`Minutes · updated register`

TERMS YOU'LL SEE

Ten terms used across the kit

Lite / Standard / Regulated	Adoption tiers. Each control says the lowest tier it applies to.	**Control ID**	A testable requirement, e.g. `AGK-STD-A03-C09` (document, then control number).
AIGC	AI Governance Committee. Can be three people or an existing risk committee.	**AI Governance Lead**	The person who runs the kit day to day.
T0–T3	AI tool tiers, from consumer tools (T0) to agents that can act (T3).	**Data labels**	Public, Internal, Confidential, Restricted, Secrets.
A0–A4	How much an agent may do on its own, from suggest-only to high autonomy.	**Connector**	A link that lets an AI tool read or act in another app, including MCP servers.
Vendor Annex	The only document naming vendor settings; one per AI platform.	**Shadow AI**	AI tools used without approval, often on personal accounts.

DOCUMENT MAP

What each tier adopts: all 48 documents

ID	Document	Owner	Lite	Standard	Regulated
FRAMEWORK					
AGK-000	Framework & Document Architecture	AI Governance Lead	● Required	● Required	● Required
AGK-001	Lite Quick Start & Checklist	AI Governance Lead	● Required	– Optional	– Optional
CORE POLICIES					
AGK-POL-01	Governance Charter & Operating Model	CEO / Board	◖ Condensed	● Required	● Required
AGK-POL-02	AI Acceptable Use Policy	HR + CISO	● Required	● Required	● Required
AGK-POL-03	AI Risk & Data Governance Policy	Risk / DPO	◖ Condensed	● Required	● Required
CORE PROCEDURES					
AGK-PRC-01	Use Case Intake & Approval	AI Governance Lead	◖ Condensed	● Required	● Required
AGK-PRC-02	Risk & Impact Assessment	Risk / Compliance	– Optional	● Required	● Required
AGK-PRC-03	Vendor & Model Assessment	Procurement + CISO	– Optional	● Required	● Required
AGK-PRC-04	AI Incident Response	CISO + DPO	◖ Condensed	● Required	● Required
AGK-PRC-05	Review & Recertification	AI Governance Lead	– Optional	● Required	● Required
AGK-PRC-06	Red-Team & Adversarial Testing	CISO	– Optional	● Required	● Required
AGK-PRC-07	AI Discovery & Shadow AI	AI Governance Lead	● Required	● Required	● Required
AGK-STD-C01	AI Literacy & Training	HR + AI Gov Lead	◖ Condensed	● Required	● Required
AGK-STD-C02	Transparency & Disclosure	Legal + Product	– Optional	● Required	● Required
AGK-GDL-C01	Platform Selection & Evaluation Guide	AI Governance Lead	– Optional	● Required	● Required
AGK-GDL-C02	Adoption Roadmap & Use-Case Prioritisation	AI Sponsor / Gov Lead	● Required	● Required	● Required
AGK-GDL-C03	Change & Communication Playbook	HR / AI Gov Lead	◖ Condensed	● Required	● Required
TRACK A · USE					
AGK-STD-A01	Platform Configuration Baseline	AI Platform Admin	◖ Condensed	● Required	● Required
AGK-STD-A02	AI Data Handling	DPO + CISO	● Required	● Required	● Required
AGK-STD-A03	Connector & Integration	CISO	◆ When switched on	● Required	● Required
AGK-STD-A04	Browser & Computer Use	CISO	◆ When switched on	● Required	● Required
AGK-STD-A05	Skills, Plugins & Instructions	AI Platform Admin	◆ When switched on	● Required	● Required
AGK-PRC-A01	Connector & Tool Request	AI Platform Admin	◆ When switched on	● Required	● Required
AGK-GDL-A01	Safe AI Use Guide for Staff	AI Governance Lead	● Required	● Required	● Required
AGK-GDL-A01-BM	Panduan Penggunaan AI Selamat (BM)	AI Governance Lead	● Required	● Required	● Required
TRACK B · BUILD					
AGK-STD-B01	Agent Development Lifecycle	Head of Engineering	– Optional	● Required	● Required

ID	Document	Owner	Lite	Standard	Regulated
AGK-STD-B02	Agent Identity & Tool Authorization	CISO + Platform	– Optional	● Required	● Required
AGK-STD-B03	Evaluation & Release Gates	Head of Engineering	– Optional	● Required	● Required
AGK-STD-B04	Guardrails & Security	CISO	– Optional	● Required	● Required
AGK-STD-B05	Monitoring, Logging & Remediation	Platform / SRE	– Optional	● Required	● Required
AGK-STD-B06	Model, Prompt & Config Lifecycle	Platform	– Optional	● Required	● Required
TEMPLATES					
AGK-TPL-01	AI System Register (Excel)	AI Governance Lead	● Required	● Required	● Required
AGK-TPL-02	Connector & Tool Registry (Excel)	AI Platform Admin	◆ When switched on	● Required	● Required
AGK-TPL-03	Use Case Intake Form	AI Governance Lead	● Required	● Required	● Required
AGK-TPL-04	Risk & Impact Assessment Form	Risk / Compliance	– Optional	● Required	● Required
AGK-TPL-05	Vendor Assessment Checklist (Excel)	Procurement	– Optional	● Required	● Required
AGK-TPL-06	Incident, Nonconformity & CAPA Log (Excel)	CISO	● Required	● Required	● Required
AGK-TPL-07	Acceptable Use Acknowledgement	HR	● Required	● Required	● Required
AGK-TPL-07-BM	Pengakuan Dasar Penggunaan AI (BM)	HR	● Required	● Required	● Required
AGK-TPL-08	Control Library, Lite Checklist & ISO SoA (Excel)	AI Governance Lead	– Optional	● Required	● Required
AGK-TPL-09	Platform Evaluation Scorecard (Excel)	AI Governance Lead	– Optional	● Required	● Required
AGK-TPL-10	Adversarial Test Catalogue (Excel)	CISO	– Optional	● Required	● Required
AGK-TPL-11	Incident Tabletop Exercise Pack	AI Governance Lead	– Optional	● Required	● Required
AGK-TPL-12	Use-Case Prioritisation & Benefits (Excel)	AI Governance Lead	● Required	● Required	● Required
VENDOR ANNEXES					
AGK-ANX-C	Claude (Anthropic)	AI Platform Admin	– Optional	– Optional	– Optional
AGK-ANX-O	OpenAI	AI Platform Admin	– Optional	– Optional	– Optional
AGK-ANX-M	Microsoft	AI Platform Admin	– Optional	– Optional	– Optional
AGK-ANX-G	Google	AI Platform Admin	– Optional	– Optional	– Optional

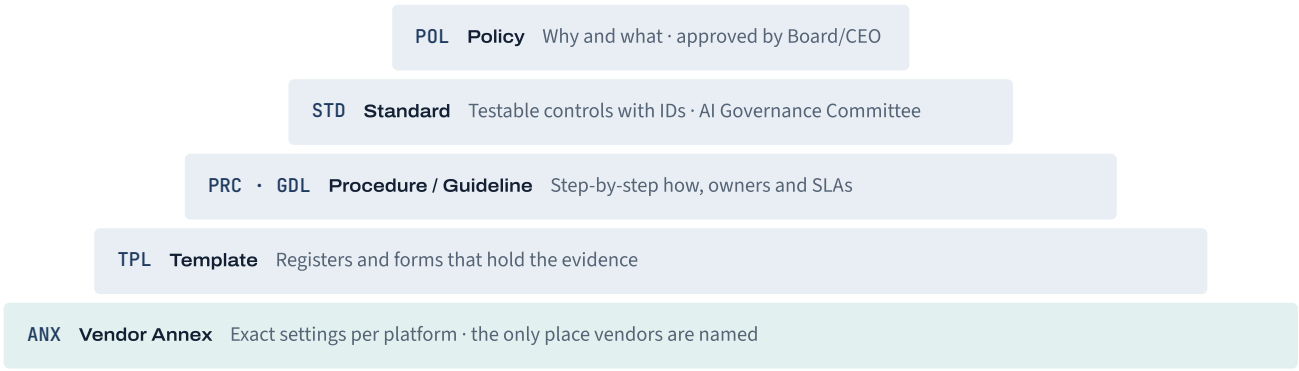
● Required ◖ Condensed version is fine ◆ Required once the feature is switched on – Optional, or only if you use that platform

PART 3 **How the kit is built**

HOW IT'S BUILT

Five levels, from board principle to platform setting

Each level answers a different question and changes at a different speed. Policies stay stable for years. Vendor settings change monthly, so they live at the bottom where updating them never needs board approval.



Follow one requirement down the levels

POLICY	AI agents shall not take irreversible actions without human approval.
STANDARD	**AGK-STD-A03-C09** Each write or act tool on a connector is set to "needs approval" or "blocked".
PROCEDURE	**AGK-PRC-A01** How a request to enable a write tool is triaged, approved and logged.
TEMPLATE	**AGK-TPL-02** Registry row records the tool, its setting and the approver.
ANNEX	**AGK-ANX-C** On Claude: per-tool permission set to "Needs approval". Each other vendor annex gives its own setting.

FRAMEWORKS BEHIND THE KIT

Built on ISO/IEC 42001, risk-managed with NIST, mapped to local rules

The kit does not invent its own framework. It follows one international management-system standard for structure, one risk framework for method, and maps every document to the rules that apply in Malaysia, Singapore and beyond.

STRUCTURE

ISO/IEC 42001:2023

The AI management system standard. Gives the kit its plan–do–check–act cycle, the policy and roles at the top, and the 38 Annex A controls. TPL-08 holds a ready Statement of Applicability.

RISK METHOD

NIST AI RMF 1.0 and GenAI Profile

Govern, Map, Measure, Manage. Shapes how risks are identified, assessed and treated in the procedures.

LOCAL RULES

Malaysia and Singapore

Malaysia AIGE principles, PDPA 2010 (as amended), draft AI Governance Bill and JPDP framework; Singapore Model AI Governance Frameworks incl. Agentic AI, and MAS AI risk guidelines for financial institutions.

SECURITY AND REACH

OWASP and EU AI Act

OWASP Top 10 for LLM and Agentic Applications for threats and testing. EU AI Act articles where outputs reach the EU.

The ISO/IEC 42001 cycle, and which documents run each part

Plan

Context, policy, risk appetite, objectives, risk and impact assessment, adoption roadmap

`POL-01 · POL-03 · PRC-02 · GDL-C02`

Do

Run AI safely: intake, vendors, discovery, training, Track A and Track B controls

`PRC-01 · PRC-03 · PRC-07 · STD-C01 · STD-A01–A05 · STD-B01–B06`

Check

Review, recertify, audit, test controls and red-team, measure benefits

`PRC-05 · PRC-06 · TPL-08 · TPL-10 · TPL-12`

Act

Handle incidents, correct nonconformities, management review, improve

`PRC-04 · TPL-06 · PRC-05 (management review)`

Where each part of the kit maps

Kit area	ISO/IEC 42001	NIST AI RMF	Malaysia AIGE / PDPA	Singapore MGF	MAS AIRG	EU AI Act	OWASP
Charter and policies (POL-01–03)	●	●	●	●	●	○	
Intake, risk, vendor, discovery (PRC-01–03, 07)	●	●	●	●	●	○	
Incidents and testing (PRC-04, 06, TPL-10, 11)	●	●	●	●	○	○	●
Review, audit, improvement (PRC-05, TPL-06, 08)	●	●	○	○	●	○	
Training, transparency, change (STD-C01–02, GDL-C03)	●	●	●	●	○	●	
Track A: using AI tools (STD-A01–A05)	●	○	●	●	○	○	●
Track B: building agents (STD-B01–B06)	●	●	○	●	●	●	●

● directly addressed · ○ partly or indirectly. Every document ends with a "Framework alignment" table, and AGK-TPL-08 maps all 422 controls individually.

How every document is laid out

All documents share the same anatomy, so a reader who knows one can navigate any of them.

1 **Document control**
ID, template version, your client version, owner, approver, review cycle, parent documents, and the tier status

2 **Template guidance**
Yellow boxes explaining how to adapt the document; deleted before issue

3 **Purpose and scope**
Why the document exists and what it covers

4 **Requirements**
Controls with IDs (e.g. AGK-STD-A03-C09), each tagged for the tiers it applies to

5 **Roles and responsibilities**
Who does what, often as a RACI table

6 **Compliance and exceptions**
How compliance is checked and how time-bound exceptions are approved

7 **Framework alignment**
The ISO, NIST, Malaysian, Singaporean, EU and OWASP references the document supports

8 **Version history and approval**
Template version history from Anchor Sprint, then your sign-off and client revision record

TWO TRACKS, ONE CORE

Using AI and building AI are governed separately

Track A is for everyone who uses AI tools. Track B is for teams that build agents and AI systems. Both share one core, so a SaaS assistant and an in-house agent sit in the same register, use the same risk tiers and follow the same incident procedure.

Track A · AI use

`all staff · admins`

How do our people use AI tools safely?

- Platform configuration baseline
- What data can go into which tool
- Connectors and MCP servers
- Browser and computer-use agents
- Skills, plugins and custom instructions
- Connector request lanes · staff safe-use guide

Track B · AI build (AgentOps)

`engineering · platform`

How do we build and run AI systems and agents safely?

- Agent development lifecycle and design gates
- Agent identity, access and tool authorisation
- Evaluation and release gates
- Guardrails and security
- Monitoring, logging and remediation
- Model, prompt and configuration change

Core · applies to both

`governance · risk · privacy`

Charter and operating model · Acceptable Use Policy · Risk & Data Policy · AI discovery · adoption roadmap · change & communication · use-case intake · risk & impact assessment · vendor assessment · incident response · red-team testing · periodic review and management review · training · transparency

SHARED RULES

One vocabulary across every document

Four definitions do most of the work. They let a staff member, an engineer and an auditor talk about the same risk in the same words.

Which data can go into which AI tool

Tool tier ↓ · Data label →	Public	Internal	Confidential	Restricted	Secrets
T0 Consumer / unapproved AI	Yes	No	No	No	Never
T1 Approved enterprise AI SSO, no training on our data, audit logs	Yes	Yes	For business need	No	Never
T2 Hardened enterprise AI + retention, transfer assessment, keys/residency	Yes	Yes	Yes	Conditional approved use case, DPIA done, minimum data	Never

Tool tier ↓ · Data label →	Public	Internal	Confidential	Restricted	Secrets
T3 Agentic AI that can act write connectors, browser, autonomous	Yes	Agent scope	Agent scope human approves write/send	As T2 autonomy up to A2 unless AIGC approves	Never

Full matrix in `AGK-STD-A02`. Secrets means passwords, API keys, tokens and private keys.

How much an agent may do on its own



Every agent declares its level at design time. Higher levels need more controls (`AGK-STD-B01`).

Use-case risk tiers

- **Prohibited**: not permitted (e.g. manipulative or social-scoring uses)
- **High**: decisions about people, material harm, A3–A4 agents → AI Governance Committee
- **Medium**: Confidential data, contained risk → AI Governance Lead
- **Low**: productivity with human review → self-service within the AUP

Data labels

- **Public**: published material
- **Internal**: everyday business information
- **Confidential**: contracts, pricing, financials, most code
- **Restricted**: personal data at scale, regulated or client-secret data
- **Secrets**: credentials, never entered into any AI tool

DAY-TO-DAY FLOWS

What the kit looks like in practice

Someone asks for a new connector

STAFF
Submits request
Business need, data involved, account type
`AGK-PRC-A01`

PLATFORM ADMIN
Triages into a lane
Green, Amber or Red within agreed SLA
`AGK-STD-A03`

APPROVER
Decides
Read-only by default. Write tools set to needs-approval
`AGK-STD-A03-C08/C09`

PLATFORM ADMIN
Configures & records
Registry row before enabling
`AGK-TPL-02`

MONTHLY
Reconciles
Enabled vs approved. Gaps become incidents
`AGK-STD-A03-C18`

Green · pre-approved
Read-only, up to Internal data, company account, vetted publisher

Amber · review
Write actions, Confidential data, new vendor or broad scopes

Red · decline or escalate
Restricted data without T2, personal accounts, unvetted publisher

A team builds an agent

INTAKE
Register the use case
Triage and risk tier
`AGK-PRC-01 · TPL-03`

DESIGN
Declare autonomy A0–A4
Threat model, identity, tool allowlist
`AGK-STD-B01 · B02`

EVALUATE
Pass release gates
Accuracy, safety, privacy, injection tests
`AGK-STD-B03`

OPERATE
Monitor & log
Daily safety scans, weekly drift review
`AGK-STD-B05`

IF IT MISBEHAVES
Remediate
Suspend → lockout → repair → reintegrate
`AGK-STD-B05 · PRC-04`

WHO DOES WHAT

Named people, not committees alone

Board / CEO

Approves the charter and AI risk appetite. Receives an annual report.

AI Governance Committee

Approves standards, high-risk use cases, A3–A4 agents and exceptions. Can be your existing risk committee.

AI Governance Lead

Runs the kit day to day: intake, registers, reviews, training.

AI System Owner

The named business owner accountable for each AI system or agent.

AI Platform Administrator

Configures platforms to the baseline. Keeps registries and annexes current.

CISO / IT Security

Security review of connectors and agents. Leads incident response.

Data Protection Officer

PDPA compliance, impact assessments, transfer assessments, breach notification.

Every employee

Uses approved tools, follows the data rules, reviews AI output, reports problems.

PART 4 **Context and support**

WHY NOW · STATUS AS AT OCTOBER 2026

AI rules in Malaysia and Singapore are moving from voluntary to expected

Most organisations already use AI tools, and many now connect them to email, files and business systems, or let agents act in a browser. That moves AI risk from "what did someone type" to "what can this system do on our behalf".

Regulators are converging on the same expectations: know what AI you run, classify its risk, name a human owner, control third parties, and report incidents. The kit gives you those building blocks on day one.

Instrument	Status (Oct 2026)	What it expects
Malaysia: National Guidelines on AI Governance & Ethics	Voluntary, in effect	Seven principles incl. accountability, transparency, privacy
Malaysia: PDPA as amended 2024	In force	DPO, 72-hour breach notification, cross-border transfer rules
Malaysia: AI Governance Bill	Draft (consultation)	Developer and deployer duties, harm tiers, incident and near-miss reporting
Malaysia: JPDP AI & Personal Data Framework	Draft (consultation)	Privacy by design and impact assessment across the AI lifecycle
Singapore: Model AI Governance Framework for Agentic AI	Voluntary, in effect	Bound agent risks upfront, human accountability, technical controls
Singapore: MAS AI Risk Management Guidelines	Final, applies from 2027	Inventory, risk materiality, lifecycle controls (financial institutions)

Issuers: Ministry of Science, Technology and Innovation (AIGE); Personal Data Protection Commissioner, JPDP (PDPA and the AI & personal data framework); Government of Malaysia public consultation (AI Governance Bill); IMDA (Model AI Governance Framework for Agentic AI); Monetary Authority of Singapore (AI Risk Management Guidelines). Status checked October 2026; confirm the current position before relying on it.

QUESTIONS LEADERS ASK

FAQ

Do we need all 48 templates?

No. A Lite adopter uses about 20 of them and 53 controls. The rest are added when you switch on connectors or browser agents, build agents, or become regulated. The two-minute tier check near the front of this pre-read tells you where you are.

How long does it take?

Lite: about 90 days to be in a stable state, with the most important steps (policy, data rules, approved tool) done in the first three weeks. Standard: three to six months depending on how much AI you build.

Is this legal compliance?

It is a governance framework mapped to Malaysian and Singaporean laws and guidelines, ISO/IEC 42001, NIST and the EU AI Act. It is not legal advice. Regulated organisations should have their pack reviewed by counsel.

Do we need ISO/IEC 42001 certification?

Not to start. The kit follows the standard's structure, so if certification becomes necessary (often because a customer asks), the Statement of Applicability in `AGK-TPL-08` is already started.

Does it depend on one AI vendor?

No. Policies and standards are vendor-neutral. Vendor Annexes cover Claude, OpenAI, Microsoft and Google; you use the ones for the platforms you run.

Will this slow AI adoption down?

It should do the opposite. Clear rules and an approved tool remove the uncertainty that stops people trying AI, and the adoption roadmap (`AGK-GDL-C02`) picks the use cases worth doing first.

What does it cost to run?

Mainly time: about half a day a week for the AI Governance Lead and an hour a month for the committee at Lite. Some controls need platform licences with admin features (single sign-on, audit logs).

WORKING WITH THE TEMPLATES

What to keep, what to change

Keep unchanged

- Document structure and IDs
- Control IDs and their intent
- The shared definitions: labels, tiers, autonomy levels

Keeping IDs lets you reuse the control library and adopt future kit updates without re-mapping.

Tailor to your organisation

- Names, roles, committee membership
- Thresholds, SLAs and review periods
- Approved tools and vendor annexes
- Sector overlays (BNM, MAS, healthcare)
- Regulator notification list

Two versions on every document

Template version

Set by Anchor Sprint, for example `AGK-POL-02 v1.0`. Shown under the title, in Document control, in the footer and in the file name, with a full change history at the end. Minor (1.0 → 1.1) for wording and defaults; major (1.x → 2.0) when a mandatory control or approval rule changes.

Client version

Yours, starting at 0.1 for a draft and 1.0 when approved. The first row of your revision history records which template version you adopted, so you can see what changed when a new template is released.

The full register of document versions is in AGK-000 (Section 4) and on the Documents & Versions sheet of AGK-TPL-08.

Reading the templates. Yellow highlighted [placeholders] are values you set. Yellow "Template guidance" boxes explain how to adapt a section and are deleted before issue. Controls marked "All tiers", "Standard + Regulated" or "Regulated only" tell you whether they apply to you.

Where to start

1. Read AGK-001
Lite Quick Start, or AGK-000 for the full framework

2. Run AGK-PRC-07
Find the AI already in use

3. Draft POL-02
The Acceptable Use Policy is the first thing staff see

4. Configure with the annex
For each AI platform you use

This kit is not legal advice. Status of draft laws reflects October 2026. Regulated-tier adopters should have their pack reviewed by Malaysian or Singaporean counsel.

NEXT STEP

From free templates to a governed AI rollout

Start free and adopt the kit yourselves. When you need more, there are three ways to go further.

Free · AI Governance Kit Lite
The 12-document Lite pack, this pre-read and the tier check. Enough to govern vendor AI tools in about 90 days.

AI Governance Hub · planned
The kit as an AI risk management platform: all 48 documents, AI system and risk register, control owners and attestation, incidents, KRI dashboard and board report, and a policy assistant. RM 2,400 per organisation per year. Waitlist open.

Deploy in your environment · planned
The same platform in your own cloud account, or fully offline on your own servers with a private AI model, for organisations that cannot use shared cloud software. Pricing on request.

Set-up and managed services
We set up owners, policy, register and an approved AI tool with you, harden your AI platforms, and keep annexes and reviews current. Quoted per engagement.

One next step. Book an online meeting or join the Hub waitlist at `anchorsprint.com/ai-governance`. Planned features may change before launch.

Anchor Sprint
Contact: `general@anchorsprint.com` · Anchor Sprint Sdn Bhd, Malaysia & Singapore