SAGAI

Strategy, Architecture, Governance and AI

Sagai Governance: build your governance with AI

A method for designing, building, testing, using and improving your governance framework. AI agents draft and prepare gate evidence. People decide.

7	10	7	7	4	15	300
stages	templates	gates	controls	tests	areas	domains

THE LIFECYCLE | seven stages, in one order, with templates for each and a decision gate that closes it

1 Assess

Assessment report: sources, findings with IDs, overlaps and gaps

Gate DG1

Existing Governance Assessment Report

2 Choose

Adopt, adapt or create for each part, with reasons. Findings answered

Gate DG2

Adopt, Adapt or Create Decision Record

3 Design

Outcomes, parts, roles, rights. Risks and impacts in six dimensions, each with its legal angle. Policy outline

Gate DG3

Governance Framework Design Brief
Governance Area Policy

4 Build

Agents draft, policy text included, under the seven controls. People edit and accept

Gate DG4, at the start

Agent Task Brief and Controls Check

5 Test

Tabletop, walk-through, evidence check, consistency check

Gate DG5

Framework Test Plan and Results

6 Put into use

Rollout. Close-out pack sent to the auditors. Policy approved for issue

Gate DG6

Rollout Record
Governance Area Policy
Findings Close-out Report

7 Improve

Review on schedule. Auditor responses recorded. Open items return to Assess

Gate DG7

Improvement and Review Log

Boundaries. A method, not answers: you give your AI agents the charter, decision rights, risk appetite, policies and assurance arrangements. The scope is to build your governance framework with your AI agents, not the lives of AI agents, which a framework such as ADLC handles. What to govern stays with you. The framework takes you through its gates, helps you agree autonomy levels and records escalations, and enforces none of them. It works by itself, and no other Sagai framework is needed. Nothing here is legal advice or governance advice.

EIGHT DECISIONS PEOPLE KEEP

1. **Which areas are in scope**
2. **How much risk is accepted**
3. **Who holds each decision**
4. **Whether policy text is approved**
5. **Who gives independent assurance**
6. **Whether an exception is granted**
7. **Whether a finding is accepted as a risk**
8. **When it goes into use**

A proposal has no standing until a person accepts or declines it.

WHAT AGENTS PREPARE

Well suited

Lists of documents and findings, matching obligations to the parts that answer them, finding gaps and overlaps, tabletop scenarios, evidence and consistency checks, first drafts and the close-out report for people to edit.

Much less so

Decision rights and risk appetite. An agent lays out options with reasons, records what was decided, and stops.

Agents never edit their own controls. The drafter is never the tester.

SEVEN CONTROLS FOR THE AGENTS IN USE

1. A named owner
2. A written task list
3. Read by default
4. Marked as draft
5. Prompts and outputs kept
6. Approved tools only
7. A person can switch it off

The minimum, not the limit. For tighter control, use a framework built for AI agents, such as ADLC.

FOUR TESTS

1. **Tabletop exercise**
   Walk a scenario through with the people involved
2. **Control walk-through**
   Follow one control from policy to evidence
3. **Evidence check**
   Does the evidence exist, current and findable?
4. **Consistency check**
   Contradictions, gaps and overlaps

Before use: a tabletop and a consistency check.

TEN CRITERIA FOR A FINISHED FRAMEWORK

Blue: shared Sagai core. Teal: added by Governance.

- Every part has an accountable role
- Effort fits the stakes
- Decisions recorded with reasons
- Reviewed on a schedule
- Roles, not people
- Clear decision rights
- Risk appetite set by leadership
- Every policy has its evidence
- Exceptions have an owner and an end date
- Checked by someone independent

GATES, AUTONOMY LEVELS, ESCALATION

7 DECISION GATES each passed when a person records a decision

DG1 | DG2 | DG3 | DG4 | DG5 | DG6 | DG7

Assess, Choose, Design, Start of Build, Test, Put into use, Improve. DG3 and DG6 need leadership by default.

3 AUTONOMY LEVELS limits, not a ladder to climb

A Advisory | B Drafting | C Checking

At every level an agent only prepares. It cannot decide, approve, publish or send, or change a document in use.

7 ESCALATION TRIGGERS The agent stops and hands the item to a person in a set role. All recorded in the Gates and Escalation Record.

AREAS, 300 REFERENCE DOMAINS, ADLC

1. Corporate, enterprise and strategic
2. Business and operational
3. Financial, investment and commercial
4. Risk, compliance, legal and assurance
5. People, workforce and capability
6. Information, data, privacy and records
7. IT, technology and digital
8. Software, engineering and AI
9. Portfolio, programme, project and product
10. Customer, market and commercial service
11. Supply chain, procurement and third party
12. Health, safety, environment and sustainability
13. Knowledge, research, innovation and IP
14. Communications, reputation and public accountability
15. Industry-specific and specialist

A reference list of 300 numbered domains under 15 area cards, copied into the organisation's own register and changed to suit. For governing AI agents themselves, use a framework built for them, such as ADLC.

FINDINGS, CLOSE-OUT AND THE AUDITORS

Raised. Each finding, overlap and gap has an ID (F-01 and so on) and an area and domain number.

Answered. Each management response names the stage and gate it is due by. Every template has a findings-addressed table.

Gate rule. A gate does not pass while a finding due at that stage has no recorded response.

Policy document | Rollout record | Close-out report

Sent to the auditors at DG6. Only the auditor who raised a finding closes it. Answers go in the Improvement and Review Log, are reviewed at DG7, and open items return to Assess. An agent never closes a finding.

 CC BY 4.0, (c) 2026 Aquasila Pty Ltd