SAGAI

Strategy, Architecture, Governance and AI

Sagai Strategy: build your strategy with AI

Work with your own AI agents to design, build, test, use and improve your strategy framework. Agents prepare drafts. People decide.

7	10	6	7	4	11	106
stages	templates	gates	controls	tests	areas	types

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

1 Assess

Assessment report: constraints, findings with IDs, overlaps and gaps. Strategy as written against strategy as funded. Register started

Gate SG1

Current Strategy Assessment Report
Objectives, Initiatives and Assumptions Register

2 Choose

Adopt, adapt or create for each part, with reasons. A decision record that says what was given up

Gate SG2

Adopt, Adapt or Create Decision Record
Strategic Decision Record

3 Design

Six dimensions, the legal angle of each, and the decisions that need leadership

Gate SG3

Strategy Framework Design Brief

4 Build

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

Gate SG4, at the start

Agent Task Brief and Controls Check

5 Test

Tabletop on a real proposal, walk-through, evidence check, consistency check. They show traceability only

Gate SG5

Framework Test Plan and Results

6 Put into use

Rollout through the organisation's own route. Register and review calendar working from day one

Gate SG6

Rollout Record

7 Improve

Review on schedule. What was reversed or dropped is recorded. Open items return to Assess

No gate of its own

Improvement and Review Log

Boundaries. A method, not answers: you give your AI agents direction statements, choices, objectives, measures, assumptions and resource rules. The scope is to build your strategy framework with your AI agents, and not plans. The choice of markets, products, investments and partners 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 strategy advice.

ELEVEN DECISIONS PEOPLE KEEP

1. **Which areas are in scope**
2. **The direction statements**
3. **Which option, and what is given up**
4. **How much risk is accepted**
5. **Which decisions need leadership**
6. **Objectives, owners and measures**
7. **How money, people and attention are split**
8. **Whether an assumption stands**
9. **When something is reversed or dropped**
10. **Whether a finding is accepted as a risk**
11. **When it goes into use**

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

WHAT AGENTS PREPARE

Well suited

Lists of what is documented, objectives and assumptions gathered from the sources, each initiative traced to an objective, measures checked for a baseline, source and date, contradictions between objectives, first drafts for people to edit.

Much less so

The choice itself. 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.

FOUR TESTS

1. **Tabletop exercise**
   Walk a real proposal through with the people involved
2. **Control walk-through**
   Follow one control from the rule to the evidence
3. **Evidence check**
   Does the evidence exist, current and findable?
4. **Consistency check**
   Initiatives to objectives, measures to their source

A clean result shows traceability, not that the direction is right.

TEN CRITERIA FOR A FINISHED FRAMEWORK

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

- Every part has an accountable role
- Effort fits the stakes
- Decisions recorded with reasons
- Reviewed on a schedule
- Roles, not people
- Choices state what was given up
- Objectives trace upward
- Assumptions are dated and watched
- Measures can be traced
- Reversals are recorded

GATES, AUTONOMY LEVELS, ESCALATION

6 DECISION GATES each passed when a person records a decision

SG1 | SG2 | SG3 | SG4 | SG5 | SG6

Assess, Choose, Design, Start of Build, Test, Put into use. SG3 and SG6 need leadership by default. Improve has no gate.

3 AUTONOMY LEVELS limits, not a ladder to climb

A Advisory | B Drafting | C Checking

At every level an agent only prepares. An agent cannot decide, approve, publish or send, or change a live document or register.

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, 106 REFERENCE TYPES

1. Direction and posture
2. Portfolio, scope and geography
3. Competition, markets and customers
4. Products, services and innovation
5. Finance and value
6. Operations, supply and delivery
7. People, culture and change
8. Technology, data and security
9. Risk, resilience and compliance
10. Stakeholders, society and partnerships
11. Performance, change portfolio, execution

Eleven cards, each type tagged by level and horizon. 28 need expert or lawyer review. Copied into the organisation's own register and changed to suit.

THE REGISTER, AND THE TRACE

Register. Objectives, measures, initiatives and assumptions. Each objective traces up, each initiative to an objective, each assumption has a source, a date, an owner and a signal. Started in Assess, checked at Test, kept current.

Given up. Every strategic choice records the options not taken and what choosing this one costs. Reversals and drops are recorded with the reason.

Works by itself. Other Sagai frameworks meet this one only at set points, and none is needed. Risk appetite and key capabilities come in as constraints the organisation states.

 CC BY 4.0, (c) 2026 Aquasila Pty Ltd