

The Institutional AI Framework

Eight operating decisions for moving from AI experimentation to a capability an institution can own, govern and fund.

Published 2026-05-22 · Updated 2026-10-02 · Version 0.1

Sources through 2026-09-30. A proposed decision method; no validated maturity score or completed field study.

Cite as: World Council for AI (2026-05-22). The Institutional AI Framework. Version 0.1. Updated 2026-10-02. worldcouncil.ai/research/institutional-ai-framework.

An AI pilot becomes an institutional capability when the organisation can operate it without depending on the pilot team's informal knowledge. Someone owns the outcome. People can challenge a result, investigate an action and recover interrupted work. The budget covers the service they are actually running.

That is a demanding transition even when the model performs well. A demonstration can establish that a task is technically possible while leaving the permissions, staff responsibilities and investment case unresolved. Senior leaders need a way to examine those decisions together.

This framework proposes eight dimensions for that examination. Use it for a particular capability and operating boundary: a reconciliation assistant in two entities, for example, rather than “our AI maturity”. It is a decision framework based on public guidance and an editorial synthesis. It is not an assessment validated against organisational outcomes, a certification or a ranking of institutions.

Executive summary

The useful question is whether the organisation has enough evidence to authorise its next operating step. The answer depends on the work, the authority granted and the consequences of failure.

Three propositions guide the framework. First, assess the complete process, including the work people retain. Second, match authority to tested capability: an assistant that proposes a match and an agent that posts an accounting entry require different release evidence. Third, test the commercial case against the same staffing and recovery plan used to approve the service. Savings that remove the people needed for fallback undermine that approval.

These are proposed design judgements. Singapore's [Model AI Governance Framework for Agentic AI](#), released on 22 January 2026, supports bounded powers, meaningful human accountability, lifecycle controls and user training. The [NIST AI Risk Management Framework](#) supports considering risk throughout development and use. Neither source validates WCAI's eight dimensions or supplies a universal threshold for release.

Read the framework with the process owner, finance sponsor, technology owner and risk colleagues. Attach evidence to each consequential decision. Where evidence is missing, record whether it prevents release, limits the scope or requires a funded test.

Why use an operating framework?

[Deloitte's January 2026 State of AI release](#) reported that 25% of respondents had moved at least 40% of their AI pilots into production. Its survey covered 3,235 business and IT leaders across 24 countries and six industries in August–September 2025. That is a self-reported global finding, not a Singapore deployment rate or proof of what caused a transition to fail.

The finding nevertheless helps separate adoption from operating capability. Buying software, conducting experiments and running a dependable business process are different achievements. A readiness discussion should examine the operating claim the organisation wants to make, rather than reward activity with a higher score.

The dimensions below are connected. Data access can constrain the architecture. The architecture can change support costs. Support costs can alter the investment case. Agent permissions can change the consequences of a mistake and the review effort required. Treating those decisions as separate workstreams can conceal a contradiction until the release meeting.

Eight dimensions, one operating decision

Dimension	Decision to resolve	Evidence to inspect
Strategy & leadership	Which outcome warrants investment, and who can change or stop it?	Baseline, sponsor, alternatives and staged funding decision
Operating model	Who performs, accepts and supports the work?	Process map, responsibilities, exception queue and operating agreement
Governance & assurance	What use and residual risk can be authorised?	Tested scope, failure evidence, approval rules and incident route
Data & knowledge	Which records may the system use, and which are authoritative?	Data flows, access tests, provenance and disputed-record process
Infrastructure & architecture	Can the service meet its obligations and withstand dependency changes?	Dependency map, service limits, recovery test and exit evidence

Dimension	Decision to resolve	Evidence to inspect
Workforce & change	Can the funded team operate the changed process?	Retained tasks, training, reviewer capacity and fallback rehearsal
Agent readiness	What actions may the system initiate or complete?	Enforced authority, action-bound approvals and duplicate/retry tests
Value realisation	What benefit survives the complete cost and workload?	Capacity, cash, quality, backlog and downside comparisons

1. Strategy & leadership

Begin with the business outcome and the decision that AI is intended to improve. A shorter reconciliation queue is more useful as an objective than “deploy agents”. Define completion, the baseline and the consequence of getting the work wrong. Consider a process change, existing software or a restricted assistant alongside a more autonomous system.

Name a sponsor with authority over the operating outcome and funding. The sponsor should be able to reduce scope or stop expenditure when evidence changes. Give each phase a question to answer and a maximum commitment. A successful technical test should lead to a specific next decision, not automatic expansion.

Attach the [investment memo](#). It makes the alternative, assumptions and continuation conditions visible. Leadership agreement is weak if it does not identify who accepts a worse-than-expected result.

2. Operating model

Map work from intake to accepted completion. Include preparation outside the application, investigation, approvals and transfers to another team. A rejected case still belongs to the operation if someone must finish it elsewhere.

Separate the business outcome owner, technical service owner and operations lead. One person can hold several responsibilities in a smaller firm; the responsibilities still need names and cover. Decide who receives an alert, who resolves an exception and who authorises a restart.

The [release and operating pack](#) provides an agreement and recovery record. Check whether the people expected to run the service can demonstrate a difficult case using those records. If only the builder can do so, the transfer remains incomplete.

3. Governance & assurance

Specify the version, users, entities, data and actions under consideration. Examine quality by consequence and by important subgroup. A useful aggregate result may conceal poor performance on the language or document type that matters most to one entity.

Agree acceptance criteria before final testing. Keep held-back cases and failure examples available to a reviewer beyond the building team. A finite test with no serious observed errors cannot establish that serious errors are impossible; the operating limits must account for uncertainty.

Use [release evidence](#) to connect results to a release decision, suspension conditions and another review after material changes. Where a use materially affects a customer, employee or counterparty, identify how the affected person can dispute evidence or an outcome, who can review it and how an accepted error is corrected. Assess applicable notification, remedy and timing obligations separately. NIST's framework is intended for voluntary use; Singapore's agentic framework is guidance. Applicable law, sector obligations and contractual responsibilities require their own assessment.

4. Data & knowledge

Identify the records the service reads and produces, where they move and who can access them. Define which source controls when records disagree. A generated summary should preserve links to the evidence that supports its material claims.

Test access by role and entity. Revoking a person's permission should also remove the access exercised on their behalf. Include knowledge maintenance: who updates the procedure when the underlying process changes, and how the organisation identifies work completed under an obsolete version.

Regional expansion requires an entity-level review of data transfers and operating conditions. The [ASEAN rollout workbook](#) helps record those decisions. A common regional vocabulary does not create one ASEAN privacy or compliance regime.

5. Infrastructure & architecture

Draw the dependencies needed to complete the work, including identity, storage, retrieval, models, tools, queues and authoritative business systems. Identify service limits and the failure of each consequential dependency.

Local deployment may give control over one layer while leaving dependence on an external model, administrator or hardware supply. An open model may permit changes while transferring maintenance and evaluation work to the institution. Compare options against the workload's actual constraints.

Inspect recovery and exit evidence. Ask whether another provider can complete representative work using an exported record, and what must be rebuilt. A contractual right to export is useful but does not establish an operationally tested substitute. Record the time, skills and cost needed to exercise the option.

6. Workforce & change

Describe the tasks people retain. Reviewers need enough evidence, time and authority to challenge a proposal. Operators need to pause processing, identify unfinished work and recover it. The service owner needs reports that distinguish more outputs from better outcomes.

Training should demonstrate those responsibilities using ordinary and difficult cases. Track checking time, corrections and queue pressure. Moving routine work into an application can concentrate harder exceptions among fewer people.

Rehearse fallback with the staffing plan used in the investment case. If the proposed savings remove contractor cover, the recovery plan cannot quietly depend on those contractors. Record peak intake, interruption and catch-up against the business deadline.

7. Agent readiness

An agent can initiate actions through connected tools. Define each permission as a business authority: which records it may read, which changes it may propose and which actions it may execute. Separate a suggestion from a commitment.

Enforce limits at the service accepting the action. An instruction inside a prompt cannot constrain credentials that retain broader authority. For consequential approvals, bind the approval to the actual amount, destination, object and action executed; changes should require another decision.

Test unexpected instructions in input, denied actions, expired permissions, repeated requests and ambiguous execution results. Financial or record-changing actions need a reliable way to inspect authoritative state before retrying. Greater autonomy needs evidence of control and recovery, beyond better generated answers.

8. Value realisation

Compare equivalent offered workloads and accepted outcomes. Count unfinished, rejected and transferred cases alongside completions. Otherwise a system can appear cheaper by passing difficult work out of the measured queue.

Separate released capacity, economic value and cash effects. Existing salaries can remain unchanged while new software and support increase spending. A capacity investment can still be worthwhile when the organisation has a credible use for the time; it should be approved on that basis.

The [economics workbook](#) provides distinct capacity, cash, downside and transition views using labelled illustrative inputs. Replace them with evidence for the proposed service. Do not treat its example as a WCAI benchmark or assume another institution will realise the same saving.

How the dimensions connect

The diagram is a proposed decision sequence, not a validated causal model. Governance, data, architecture, people and agent authority are examined together before authorising an operating boundary. Observation after release can change any earlier decision.

```

Strategy & leadership: outcome, alternatives, sponsor
      ↓
Operating model: accepted work, ownership, support
      ↓
Governance · Data · Architecture · Workforce · Agent readiness
      ↓
Authorise a bounded operation, restrict it, or hold
      ↓
Value realisation: outcomes, workload, cost, residual risk
      ↺
Revisit scope, funding and controls when evidence changes

```

Use an evidence record instead of a total score

For each dimension, record the claim, evidence location, accountable owner and unresolved decision. Use the statuses **supported for this scope**, **restricted**, **unresolved** and **not applicable with reason**. These describe the decision; they are not numbers to average.

Record	Example to replace with your evidence
Capability and boundary	Invoice match proposals for named entities; no payment authority
Decision requested	Limited live use by named reviewers
Claim	Access cannot reach another entity's records
Evidence	Allowed and denied access tests for the release version
Owner	Named technical owner; acceptance by the business owner
Status and restriction	Supported for tested users; additional entity access unresolved
Next action	Test additional role/entity combinations before expansion

This is an illustrative record. An unsupported consequential permission should not be offset by strong training or an attractive return estimate. Some constraints prevent the requested release; others justify a narrower one. The decision owner must distinguish them.

An initial working session can produce a useful result even when it authorises nothing. Choose a single capability, identify the missing evidence and fund the smallest test that resolves the next decision. That produces a more inspectable plan than an enterprise score built from incompatible self-ratings.

Method, scope and limitations

Version 0.1 combines public desk research with an editorial synthesis of operating decisions. The eight dimensions are an organising proposal, selected to connect institutional strategy with process ownership, technical authority and commercial evidence. The review prioritised original public guidance and survey releases where their scope could be identified. Sources were inspected through 30 September 2026.

This is a selective review, not a systematic literature review. Singapore's January 2026 agentic framework supplies a regional reference for delegated authority and accountability; NIST's AI RMF supplies a broader voluntary lifecycle vocabulary. Deloitte's survey contextualises pilot-to-production activity and was not used to validate the dimensions. The practical worksheets are proposed working methods and contain labelled illustrations. No client outcomes, interviews, representative organisational sample or predictive validation underpin this version.

The revision rule for this framework is to create a new documented version when the dimensions or recommendations materially change. Source updates retain their own cutoff and change note. Independent editorial and source challenges inform this edition; they are distinct from practitioner testing or human subject-matter endorsement. A released revision should identify what changed and the evidence for that change, so readers can distinguish a source refresh from a changed organising judgement.

The dimensions overlap deliberately. Agent readiness makes delegated authority explicit, while governance examines whether that authority is acceptable. Workforce and operating-model evidence may share records. Use one record where it answers both decisions rather than duplicate paperwork. For a low-consequence assistant, the review can be brief; irreversible transactions or sensitive services demand more evidence.

This framework may miss sector-specific constraints and cannot determine legal compliance. Singapore is the principal regional reference point; application in another jurisdiction requires local analysis. A future revision should test whether teams interpret the questions consistently, whether the record changes actual decisions and which dimensions fail to reveal important constraints. Until then, claims about validated maturity or comparative institutional performance would exceed the method.

Sources and working resources

- [NIST AI Risk Management Framework overview](#): voluntary lifecycle risk guidance; original framework released in 2023.
- [Singapore Model AI Governance Framework for Agentic AI announcement](#): 22 January 2026; bounds, accountability, controls and training.
- [Deloitte State of AI release](#): January 2026 release, August–September 2025 global survey fieldwork.
- [Workflow discovery pack](#), [release and operating pack](#), [economics workbook](#) and [investment memo](#): proposed practical records, with examples identified where used.