




Sovereign AI

The operating playbook

SEPTEMBER 12, 2026



Control what matters.
Keep the ability to change.



Contents

From a control requirement to a service you can trust.

A casework assistant runs through the report, connecting the design choices to their operating tests.

THE CASE FOR CONTROL

Sovereignty as an operating capability	03
What the market is buying	04

DESIGN THE SYSTEM

Define the workload	05
Choose the deployment	06
Trace data custody	07
Retain institutional knowledge	08
Constrain action	09
Rehearse continuity	10

PROVE IT IN OPERATION

Price accepted outcomes	11
Procure enforceable controls	12
Run the ninety-day program	13
Build national capability	14
Apply the four acceptance tests	15



Sovereignty is a capability, not a location.

An institution is sovereign in practice when it can keep an important AI service useful, governable and replaceable as conditions change. Local infrastructure can help. It does not, by itself, establish control over information, authority to act or the ability to leave a supplier.

The operating test

Five rights turn an aspiration into requirements. Each should be specified for a workload and demonstrated before production.

RETAIN VALUE

Keep context, evaluations and knowledge created through use.

CHANGE

Replace models, infrastructure and the control platform.

LOCATE


Choose where information is processed, stored and supported.

EXCLUDE

Control who can read records, use keys and change policy.

OPERATE

Maintain an accepted service through a defined disruption.



The playbook's central choice

Protect the institutional core while keeping access to useful external capability. Add isolation, dedicated capacity or local execution when a specific requirement justifies the burden. Avoid making every workload pay for the most restrictive architecture.

What we retain, and what we qualify

Model choice and owned context are durable principles. Hardware ownership, zero retention, open weights and reversibility require more precise tests. An export is not necessarily a working exit; a branch cannot undo a disclosure.



Start here: name the service owner, the unacceptable failure and the evidence that would demonstrate control.



A growing market. Different problems to solve.

Sovereign AI combines infrastructure, operational assurance, and models and applications. These spending pools overlap. The useful commercial question is which control requirement creates demand for a production service, not how many investment announcements can be added together.

Sovereign cloud infrastructure spending



The 2026 forecast represents 35.6% growth. This infrastructure measure includes non-AI workloads and is not total sovereign AI revenue.



Capacity and access

The EU's July 2026 call targets up to seven AI Gigafactories, with up to EUR10bn publicly funded to attract at least EUR20bn privately. India emphasizes shared access: more than 45,000 GPUs as of June and 237 projects using subsidized compute by August.



Capability and applications

The Mistral-HUMAIN collaboration links Saudi infrastructure with model localization and regulated applications. The UK's GBP500m initiative combines capital, compute and talent. Each approach still needs customers, operators and a credible delivery model.



Where the playbook creates value

Our view: assurance, deployment and workflow integration can be durable sources of differentiation. A regional operator needs utilization; a regulated buyer needs custody and continuity. Validate the funded workload and the willingness to pay before committing capacity.

An announcement establishes intent. Usable service establishes opportunity.



Write the charter before choosing the stack.

Begin with one consequential workflow. A charter connects the business objective to the controls the architecture must provide. It also prevents a supplier's product description from becoming the institution's definition of sovereignty.

WORKED EXAMPLE: CASEWORK ASSISTANT



Mission

Prepare an evidence-based case summary for a trained reviewer.



Information

Approved case records, with access tied to the reviewer and the case.



Authority

Retrieve and draft. No independent case decision or external communication.



Quality

The service owner defines acceptable accuracy, completeness and correction effort.



Continuity

Search and manual processing remain available when the model route fails.

The decisions behind the form

Separate what the assistant may recommend from what it may execute. Identify which data cannot leave an approved boundary and which service loss the business can tolerate. Translate broad terms such as confidentiality into an observable requirement.



Data custody

Keep approved case records within the defined boundary, with access tied to the reviewer and the case.



Human review

The assistant may recommend and draft. A trained reviewer makes the decision.

Assign the owners



Business owner accepts task quality and degraded service.



Security defines access and custody boundaries.



Platform team implements controls.



Procurement secures corresponding obligations.



Operations rehearses recovery.



Set the baseline

Measure current effort and cost per accepted summary. Assemble representative cases, including difficult examples and denied-access scenarios. These become the evaluation set used for model choice, rollout and future changes.

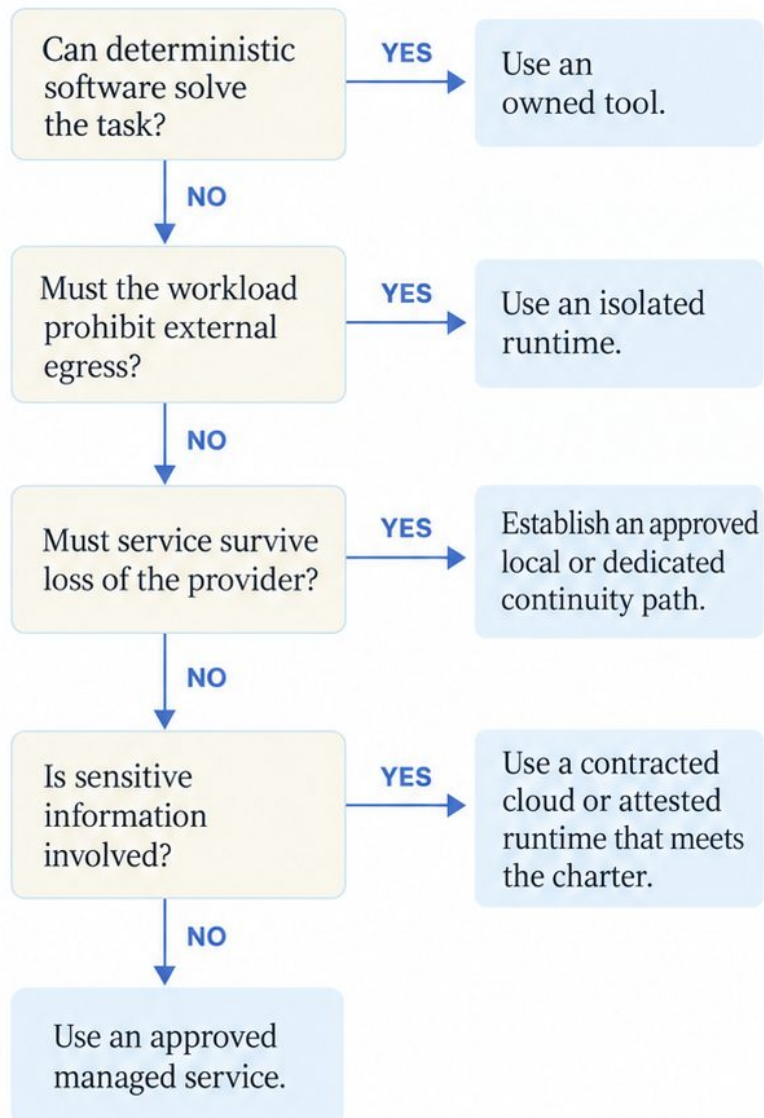
Exit this stage with an approved charter, a dependency register and a named owner for each critical control.



Route the workload, then select the model.

Architecture follows the charter. Provider nationality, data residency and effective control are distinct. Evaluate administration, model access, updates and recovery as part of the service, including dependencies that remain outside the chosen region.

A deployment decision path



Qualify the chosen route

Compare models on the institution's own cases: correctness, completion, permission violations, latency and correction effort. Open weights can enable local execution, but do not necessarily disclose training data or remove license conditions. Managed models reduce operating burden while introducing service dependencies.


Apply it to casework

Keep records, identities and permissions in the approved environment. Route only permitted tasks and information. If no approved model meets the threshold, return work to a person. A fallback must not silently weaken the original data or action constraints.

Follow every copy of the information.

A case summary can expose information through more than the model request. Prompts, retrieved records, outputs, traces and support artifacts may each have different retention and access rules. Custody is the complete path, not only the database where the record began.

Trace the casework interaction



For each copy, record location, purpose, access, retention, deletion and exceptional disclosure. Include caches, backups, support and subprocessors. Customer-controlled keys do not always mean a provider cannot process plaintext while performing its service.

Ask four different questions

Is content stored? Can it be used for training? What safety signals are retained? Who can inspect an exception? A no-training commitment does not eliminate logs or support access. Retention alone does not establish training use.

Treat new safeguards as conditional

By September 2026, providers were developing customer-controlled frontier monitoring. OpenAI's Private Safety Processing was previewed in August; Anthropic announced Enterprise Frontier Safeguards in September with phased fall availability. Confirm eligibility and the configuration actually operating today.


Test before exposure

Use identifiable synthetic records to exercise storage, deletion and revocation. Require supplier evidence where direct inspection is unavailable. An unexplained path for protected information is a reason to restrict the route.


Own what the institution learns.

The durable asset is the knowledge created through work: records, decisions, corrections and criteria for success. Keep it in a form that can outlast both the model and the platform that first captured it. An archive of conversations is not enough.

The portable core



A controlled improvement loop



Treat feedback as evidence to curate, not automatically trusted training material. Preserve evaluation and release versions. A graph or ontology can help represent relationships, but its value depends on understandable definitions and a usable export path.

Rehearse the move

Export one casework workflow into a second environment. Rebuild retrieval, restore permissions and run the same cases. Measure missing fields, manual conversion effort and quality loss. Fix those obstacles while the architecture is still flexible.

What counts as proof

The export must preserve behavior as well as content. Verify revoked access, inherited restrictions and combinations of records. Repeat after major changes to the data model or orchestration layer.


A working restore with acceptable quality is stronger evidence of independence than an export button.



Put authority outside the model.

A model may propose an action; the surrounding system must decide whether that action is allowed. Prompt instructions are useful guidance, but should not be the only mechanism preventing a disclosure, payment, record change or external message.

The action boundary



For the casework assistant



The agent may retrieve permitted evidence and prepare a draft.

It cannot send the summary or decide the case on its own.

The reviewer's approval does not grant open-ended permission for whatever the agent attempts next.

Assume retrieved material can mislead



A document or email may contain instructions that conflict with the task. Treat it as data.

Test attempts to redirect the agent, retrieve restricted records or misuse tools. The decisive question is whether the system blocks unauthorized action even when the model is persuaded to attempt it.

Control the irreversible step



Branches and version history help recover controlled software and data. They cannot recall a disclosure or every external action. Place approval before irreversible steps, scope credentials narrowly and limit outbound destinations.

Prove the limit




Exercise access revocation and the kill switch. Confirm that active work stops without leaving untracked partial actions.

Record both blocked attempts and legitimate tasks that still succeed.

An alternative must survive the rehearsal.

Owning hardware or naming a second provider does not establish continuity. The alternative needs capacity, compatible software, restored context, valid credentials and an operating team that can take over within the required interval.



Define the loss the business accepts

Specify tolerable interruption, acceptable state loss and the functions that must continue. Agree the degraded mode before an incident. Otherwise, the institution discovers during an outage that its fallback does not support the mission.

Run a timed switch

Have the on-call team follow the recovery procedure. Restore the export package, verify permissions and run representative cases. Record manual repairs, achieved quality and elapsed time. An architect's memory is not a substitute for a usable runbook.

Test shared dependencies

Exercise loss of identity, keys, connectivity, model access and the control platform. Two suppliers may depend on the same critical service. Include expired credentials and unavailable update or package services in the scenarios.

Keep the option alive

Assign responsibility for spare capacity, compatibility and repeat exercises. Rehearse after material changes. A dormant fallback that is never maintained can become another liability.



Price successful work, including control.

Token prices and GPU-hour rates are inputs. Compare the cost of accepted outcomes at an agreed quality, latency and recovery level. Include failed attempts, retries and human correction in the cost base.

$$\text{Cost per accepted outcome} = \frac{\text{Annual fixed cost}}{\text{Annual accepted outcomes}} + \text{Variable cost per accepted outcome}$$

A utilization example

$$\text{\$1.2m annual fixed cost} / 600,000 \text{ accepted outcomes} = \text{\$2 fixed cost per outcome}$$


$$\text{\$1.2m annual fixed cost} / 300,000 \text{ accepted outcomes} = \text{\$4 fixed cost per outcome}$$


Illustrative arithmetic, not a market benchmark. The same capacity can become much more expensive per useful result when demand falls.

Include the full burden

Count hardware or reservations, power, networking, redundancy, storage and deployment work. Add staffing, patching, security, evaluations, licensing, support and recovery exercises. Allocate shared costs consistently across alternatives.

Stage the build decision

Dedicated capacity is more compelling when demand is sustained, operating skills exist or continuity has strategic value. Managed services can be more efficient for volatile demand. Model low, expected and high demand rather than relying on one utilization assumption.



For casework, the denominator is an accepted summary, not a generated response. Measure reviewer effort and the cases returned for correction.



If resilience justifies a more expensive option, fund that requirement explicitly instead of hiding it inside optimistic unit economics.

Make the contract match the operating test.

Buy demonstrable controls for a specified service. The same charter should govern architecture, supplier selection and production acceptance. A broad sovereignty label cannot substitute for a contract the technical and operating teams can execute.

Requirement	Obligation	Demonstration
 CUSTODY	Defined retention and access terms.	Trace copies and exercise deletion.
 EXIT	Usable exports and transition support.	Restore a workflow elsewhere.
 AGENCY	Scoped tools and approval controls.	Attempt unauthorized actions.
 CONTINUITY	Recovery duties and support.	Complete a timed rehearsal.

Resolve the service boundary

Name the contracting entity, deployment route, hosting locations, administrators and subprocessors. Separate storage, training use, safety monitoring and support access. Clarify customer responsibilities and how material changes trigger reassessment.

Secure a working exit

Include context, permissions, evaluations and configuration, not only raw records. Specify formats, transition assistance, relevant charges and deletion confirmation. Test a representative export during evaluation. An exit clause that cannot be executed is weak protection.

Tie commitments to evidence


Use the same cases and criteria across finalists. Certifications can support assessment but do not prove the customer configuration. Where practical, stage rollout and capacity commitments against accepted milestones.

The service owner needs a clear basis to withhold production approval when an essential test fails. Procurement, security and operations should agree that basis before signature.



Ninety days to a defensible decision.

A ninety-day program should prove an operating discipline around a few important workflows. It is not a promise to complete an entire platform. Scope the pilot so the business owner can choose to release, restrict or stop on evidence.



The handover matters

Assign responsibility for monitoring, updates, keys and repeat exercises.
Keep unresolved dependencies visible with owners and consequences.
Reassess when the model, access route, data class or business use changes.


A successful pilot authorizes the next controlled step.

It does not automatically authorize more sensitive data or greater action authority.



Build national capability from usable demand.

A national program creates value when institutions turn compute into sustained services, domestic skills and useful applications. Installed GPUs are an input. The operating chain must connect demand, access, capability and adoption.



Fund the connections

Pool procurement where infrastructure and assurance can be shared, while preserving accountability for each application.

Distinguish research access, public-service deployment and commercial industry support; they may need different funding and service models.

Measure use, not only supply

Track time to obtain access, useful utilization, production applications and service quality.

Budget for networking, power, security and renewal.

Subsidized capacity that is difficult to access or unsupported can remain disconnected from demand.

Preserve future choice

Require usable exports, transition support and control over institutional context.

Diversify where it improves continuity or bargaining power, accounting for the operating burden. Do not multiply suppliers merely to appear independent.

Finance frontier training when mission, talent, data and sustained demand justify it. Otherwise, adaptation, evaluation and application delivery may create more near-term value.


Funding should follow delivered capability and measured adoption.

Four proofs before production.

The report's argument returns to a simple standard: the institution should be able to demonstrate the control it claims. Use four tests before release, after material changes and during periodic reassessment. Keep the results with the workload charter.

EXIT

Restore the workflow in a second approved environment. Preserve context, permissions and acceptable quality. An export file or a changed endpoint alone is insufficient.




CUSTODY

Verify locations, key authority, retention, deletion and exceptional access. Resolve unexplained paths for protected information before widening exposure.




AGENCY

Attempt unauthorized retrieval and action. Exercise approvals, revocation and the kill switch. Show that the surrounding system enforces the limits even if the model attempts to cross them.



CONTINUITY

Remove a critical dependency and recover an accepted service within the agreed interval. Record the achieved quality, manual repairs and remaining gaps.



Make a decision, not a presentation

For the casework assistant, approval requires an acceptable summary, a defensible data boundary, constrained action authority and a rehearsed manual fallback. If a test fails, repair the control, narrow the scope or delay release.

For suppliers, repeatable proof of these outcomes is a stronger proposition than an unqualified sovereignty promise.

Control what matters. Retain what you learn. Keep the ability to change.