

How to Build an Enterprise AI Operating Model



A practical guide to ownership, delivery and measurable value for Saudi organizations

Ownership. Decision rights. Shared services.
A repeatable path from priority workflow to business value.

September 2026 | English edition

Make AI a managed business capability

The leadership challenge: connect ambition to everyday decisions



An AI strategy can identify attractive opportunities and still leave the organization unable to deliver them. A pilot may have a sponsor but no production owner. A model may pass a technical test while the business process remains unchanged. A shared platform may be available, yet every department negotiates data access and approvals from the beginning. These are operating-model problems.

An enterprise AI operating model defines how the organization selects work, assigns authority, funds capabilities, delivers products and manages them after launch. It connects business leadership, domain teams, technology operations and control functions through explicit working agreements. Its effectiveness becomes visible when a routine decision can be made without restarting a debate about who owns it.

The execution gap is visible in external research. Deloitte's 2026 survey of 3,235 business and IT leaders across 24 countries found that 30% of surveyed organizations were redesigning key processes around AI, while 37% reported little or no process change. These are global survey findings, not a Saudi adoption estimate. [1]

Start with the workflow you intend to change

SOHOB recommends selecting one consequential workflow before designing committees or job titles. Map its current handoffs, baseline performance, users and failure consequences. Then specify what AI may recommend, what it may execute and which decisions remain with people. A logistics exception process, for example, needs ownership of shipment outcomes, escalation handling and data quality alongside model development.

SOHOB POINT OF VIEW

The operating model is ready to test when a named team can take one priority workflow from an approved outcome to a supported service, with clear authority at every handoff.

This article extends the operating-model and talent domain of SOHOB's Enterprise AI Strategy Framework. The design recommendations that follow connect organizational structure to decision rights, delivery evidence and business value. The accompanying tables are practical starting points to adapt to the organization's scale and risk profile.

Choose where decisions should sit

Centralize shared commitments; place delivery close to the business



Choosing a structure begins with two questions: how much must be consistent across the enterprise, and how much domain knowledge is needed to change the work? The answer may vary between platform engineering, product delivery and risk oversight. A single organizational label should not determine every responsibility.

Structure	When it fits	Failure mode to address
Centralized	Early capability building; scarce specialists; a small portfolio.	The central team becomes a queue for every business request.
Federated	Several capable domains need common platforms and controls.	Local teams bypass shared standards or duplicate investment.
Decentralized	Highly autonomous units with mature delivery and oversight.	Inconsistent controls, disconnected data and repeated costs.

For a multi-domain organization building toward scale, SOHOB recommends testing a federated arrangement: a central AI office coordinates standards and the portfolio, platform teams provide reusable services, and business product teams own workflow change and results. Smaller organizations can assign these responsibilities to existing roles rather than create new departments.



Shared evidence connects the portfolio, product decisions and operational outcomes.

Make the boundaries contractual in practice. The center publishes available services, onboarding requirements, response targets and exception routes. Each domain names an accountable product owner and commits staff time to adoption. Risk and assurance functions retain the authority to challenge decisions through existing enterprise governance.

DECISION RULE

Choose the structure that removes the next delivery constraint. Reassess it as domain capability grows; do not make a permanent central bottleneck the price of consistency.

Assign decision rights before delivery

Every material decision needs an accountable owner and an evidence trail



A responsibility chart is useful only when it resolves real choices. Teams need to know who can commit funding, approve data use, accept residual risk, release a product and stop it. Assign one accountable role to each decision, name the required reviewers and document who acts when that role is unavailable.

NIST's AI Risk Management Framework Playbook calls for clear responsibilities and communication, and places responsibility for AI development and deployment risk decisions with executive leadership. It also discusses separating development from testing to support independent challenge. [2] SOHOB translates those principles into the following operating decisions.

Decision	Accountable role	Evidence required
Prioritize and fund	Executive portfolio owner	Outcome baseline, investment case, capacity and dependencies.
Authorize data use	Data owner	Approved purpose, access scope and required privacy review.
Approve release	Designated release authority	Evaluation results, completed controls and service readiness.
Accept residual risk	Authorized risk owner	Documented exposure, mitigations, expiry and escalation.
Suspend the service	Service owner / incident lead	Trigger breached, containment action and recovery plan.
Validate benefits	Business outcome owner	Measured result with finance or PMO verification.

The AI office coordinates this process; it does not automatically own every approval. Embed AI decisions in existing investment, architecture and risk forums where they fit. Set review response targets appropriate to the risk tier, and escalate overdue decisions through a named executive route.

For Saudi organizations, SDAIA's AI Ethics Principles provide a national reference for accountability, privacy, safety and other ethical considerations. [3] Translate the relevant principles into assigned actions and evidence. Have the responsible privacy, legal and security functions determine applicable requirements for each workflow and data flow.

EXECUTIVE TEST

Ask who can pause an AI service today, who must be notified and who can authorize restart. If the answers conflict, resolve the operating model before expanding autonomy.

Build shared services teams will use

Reusable foundations need owners, service commitments and an adoption path



A common platform creates value when a product team can consume it reliably. Publishing technical standards without usable services simply moves implementation work into every domain. Build the shared foundation around recurring needs in the approved portfolio, then define the interface between platform providers and product teams.

Shared service	What the domain receives	Operational responsibility
Governed data and knowledge	Approved sources, access patterns and retrieval components.	Data owners maintain quality, permissions and update commitments.
Model access and evaluation	Approved model choices, evaluation tools and usage limits.	Platform team controls versions, access, availability and cost telemetry.
Delivery and monitoring	Deployment pipelines, logs, alerts and rollback mechanisms.	Platform team operates tooling; product owner sets use-case thresholds.
Enablement and assurance	Templates, specialist support and scheduled control reviews.	AI office coordinates support; control functions own their judgments.

Define the service agreement

For each reusable capability, publish its owner, supported use, onboarding steps, response expectations, charging approach and change policy. A domain should understand what happens when a data source becomes stale, a model version changes or a shared service fails. Supplier-managed capabilities need the same internal ownership and escalation clarity.

Use approved patterns as the default, with a documented exception route. Evaluate requests for a different model or platform against business need, integration effort, risk and total operating cost. Feed recurring exceptions back into the platform backlog so that standards improve rather than become detached from delivery.

Keep the product team complete

Each priority workflow needs business product ownership, technical delivery, data expertise, operational support and change capacity. Roles can be shared across products, but their time must be allocated. A part-time sponsor cannot substitute for a product owner with the authority to change process rules and resolve daily trade-offs.

OUTPUT

A short service catalogue and two-way working agreement: what the platform provides, what the domain must supply, and how either side escalates a blocked dependency.

Make the path to production repeatable

Connect SOHOB's governance gates to concrete release and support decisions



A pilot should enter a defined delivery path as soon as its outcome and scope are approved. Use the same product record from initial selection through retirement so that ownership, data permissions, evaluations and operating decisions remain connected. Reuse the four governance gates from the SOHOB strategy framework, with evidence proportionate to impact and autonomy.

Gate	Required decision	Minimum evidence
G1 Business & risk	Should this workflow enter delivery?	Named owner; baseline; value hypothesis; risk class; funding.
G2 Data & privacy	May the solution use the proposed context?	Sources; permissions; intended purpose; relevant privacy and security review.
G3 Model / agent validation	Is behavior acceptable within the approved scope?	Task tests; failure analysis; adversarial checks; human-control tests.
G4 Production & monitoring	Can the organization operate and recover the service?	Support owner; alerts; rollback; incident route; user readiness.

Prepare evidence while building. Agree evaluation criteria with domain experts before testing, and include cases where the system should refuse, defer or escalate. For Arabic-facing workflows, include representative Arabic content, terminology and user interactions in the evaluation set. Define who approves material changes to models, prompts, tools or knowledge sources after release.

Plan the difficult day before the launch day

A successful release includes an incident rehearsal. Simulate a stale knowledge source, an incorrect recommendation or an unavailable dependency. Confirm that an alert reaches the right person, the service can fall back safely and users understand the temporary process. Record the decision to resume normal operation.


Where agents can perform actions, define tool permissions, transaction limits and approval boundaries explicitly. Expand authority only after evaluation and operational evidence justify it. A change that increases impact or autonomy should trigger renewed review, even when the underlying model is unchanged.

RELEASE RULE

Production approval requires an operating owner and a recovery route. A working demonstration alone is insufficient evidence of service readiness.

Redesign work around human judgment

Specify what people do when AI succeeds, fails or needs help



Adoption becomes operational when the new workflow is easier to follow than the old one and employees know how to handle exceptions. Define the user's responsibilities alongside the system's permissions: what to verify, what to approve, what to override and when to escalate. Build training around those tasks.

The Enterprise AI Playbook supplied for this article examines 51 successful deployments. Its discussion of sponsorship and oversight highlights the organizational work surrounding delivery, while noting that appropriate oversight depends on task risk and context. Because it studies successful cases, its findings should not be treated as a universal success rate or a guaranteed productivity outcome. [4]

Illustrative workflow: customer-service case resolution

Consider a service team introducing an assistant that retrieves approved policy and drafts a response. The team leader owns the service outcome. The agent verifies the evidence and sends the response. A policy owner maintains the knowledge source. Complex or ambiguous cases go to a specialist; the assistant has no authority to approve a refund unless that action is separately authorized. This is an illustrative design, not a reported client result.

Work moment	Human responsibility	Operating evidence
Before use	Confirm case scope and permitted data.	Role-based access and completed task training.
During use	Check policy support; correct or reject the draft.	Review samples, override reasons and escalation records.
After use	Resolve customer outcome and report recurring issues.	Quality results, repeat contacts and improvement backlog.

Give managers time and incentives to support the change. Update standard operating procedures, allocate specialist coverage and make it safe to report errors. Track whether employees use the assistant in the intended workflow, not merely whether they log in. Repeated workarounds are evidence that the process or product needs attention.

ADOPTION RULE

Before rollout, every affected role should be able to explain its new task, its decision boundary and where to get help. Training completion is one input, not proof of adoption.

Fund the service and verify its value

Connect investment decisions to sustained performance



AI funding should cover the complete service: discovery, data preparation, integration, evaluation, user enablement, production support and eventual retirement. A budget that ends at launch leaves the business dependent on an unmaintained capability. Separate shared platform investment from domain product budgets, while keeping both visible in portfolio decisions.

Release funding against evidence

SOHOB recommends staged commitments. Fund discovery to establish feasibility and a credible baseline. Fund delivery when ownership, data access and an evaluation plan are agreed. Authorize expansion when the workflow is operating within its quality and risk boundaries and evidence supports the value case. Stop or redesign initiatives whose assumptions no longer hold.

Measure	How to use it	Accountable owner
Business outcome	Track cycle time, service quality, loss reduction or revenue against baseline.	Business outcome owner
Workflow adoption	Measure intended task coverage and sustained use, with reasons for bypass.	Domain product owner
Service performance	Track failures, latency, incidents and recovery against commitments.	Service owner
Unit economics	Calculate total service cost per successfully completed business task.	Product owner + finance
Control effectiveness	Review exceptions, unauthorized actions and overdue corrective work.	Relevant control owner

Time released is not automatically a cash saving. Identify whether it creates additional capacity, reduces overtime, avoids hiring or supports a verified budget reduction. Include the cost of human review, rework and operating the AI service. Use comparable workload periods or controlled comparisons where practical, and record other changes that could explain the result.

Run a weekly product review for service issues and improvement priorities, a monthly portfolio review for value and funding, and an executive review for material strategic or risk decisions. Adapt the cadence to the business; every forum should produce decisions, owners and due dates.

VALUE RULE

Expand the product when observed workflow outcomes justify its full operating cost and risk. Report the strength of the evidence alongside the benefit claim.

Put the operating model to work in 90 days

Test the organization through a live workflow, then refine the design



The first 90 days should establish a working management system and test it on a bounded workflow. Use the period to expose missing ownership, slow decisions and unclear handoffs. Production timing must follow readiness and required approvals; a calendar target should not override an unresolved control gap.

Period	Priority actions	Exit evidence
Days 1-30 Define	Choose an executive sponsor and workflow owner; map the current process; assess maturity; document baseline, risk and decision rights.	Approved operating charter; selected workflow and named decision owners.
Days 31-60 Design and prepare	Agree team capacity, platform services, data access, funding and governance gates; design the future workflow and evaluations.	Staffed delivery plan, service agreements and accepted evaluation criteria.
Days 61-90 Test and mobilize	Run a controlled release when ready; rehearse incidents; train users; review outcomes and revise the model.	Readiness decision, initial performance evidence and next-wave priorities.

Five outputs leadership should be able to use

- An operating charter describing scope, ownership and the chosen structure.
- A decision-rights matrix with escalation and exception routes.
- A shared-service catalogue with named service and data owners.
- A product record linking funding, evaluations, approvals and support.
- A value scorecard and a prioritized roadmap for the next 12 months.

Use lessons from the first workflow to improve the next. If approvals repeatedly stall, simplify the evidence request or clarify authority. If users bypass the product, inspect process fit and quality. If every domain rebuilds the same component, promote it into a supported shared service. These feedback loops are how the operating model matures.

HOW SOHOB CAN SUPPORT

Start with an Enterprise AI Operating Model assessment: clarify current ownership and delivery constraints, define the target model, and build a practical roadmap connecting platform capabilities, governance and business outcomes. Contact SOHOB to discuss the first workflow and the decisions needed to move it forward.

Evidence & further reading

Sources support the recommendations; the operating design is developed for SOHOB.



The recommendations, working agreements and decision tools in this article express SOHOB's proposed approach. External publications are cited for supporting evidence and context; their frameworks and research remain attributable to their authors.

[1] Deloitte (2026). State of AI in the Enterprise: survey findings and methodology.

Used for the global process-redesign evidence in the opening section. Survey fieldwork: August–September 2025.

[Read the source](#)

[2] NIST. AI RMF Playbook, GOVERN 2.1 and GOVERN 2.3.

Used for documented roles, executive accountability and independent challenge. NIST guidance is not presented as a Saudi legal requirement.

[Read the source](#)

[3] SDAIA (2023). AI Ethics Principles, particularly Accountability & Responsibility.

Used as a Saudi ethical reference; application-specific obligations remain a matter for the responsible functions.

[Read the source](#)

[4] Pereira, E.; Graylin, A. W.; Brynjolfsson, E. (April 2026). The Enterprise AI Playbook: Lessons from 51 Successful Deployments. Stanford Digital Economy Lab.

User-supplied reference, especially pp. 29, 36 and 105. Used for qualitative context on oversight, sponsorship and implementation; no case-study productivity figures are generalized.

[Read the source](#)

[5] Deloitte. How to Create an Effective AI Strategy. State of AI in the Enterprise, 4th Edition.

User-supplied reference, pp. 2–4. Background reading for the connection between business strategy and AI investment; older survey statistics are not used as current evidence.

[Read the source](#)

CONTINUE THE CONVERSATION

Which workflow is ready for a clear owner; an agreed operating model and a measurable outcome? Start there with SOHOB.

Web sources checked September 2026. User-supplied reports are identified above. Photographs are illustrative scenes, not documented client engagements.

[Contact Sohob Arabia](#)