FRACTO

Playbook

2026 AI Org Design & CoE Playbook

A Practical Playbook for Leadership Planning & Execution

Operating Models | CoE Charter | Role Library | 12-Month Roadmap | RACI | KPIs

What's Inside

- Operating model framework — centralized, federated, embedded
- Complete CoE charter, governance bodies, and funding models
- Role library with JD templates and stage-appropriate hiring
- Industry-specific org charts (FS, healthcare, retail, manufacturing)
- 12-month phased roadmap with weekly milestones and RACI
- Executive KPI dashboard — strategy, delivery, governance, talent
- 30-point AI org readiness checklist
- Templates: charter, intake, JD, business case, vendor scorecard

2026 Edition | Enterprise AI Series
fracto.ie • hello@fracto.ie

Executive Summary

Mostenterprises today are not limited byAI tools — they are limited by AI organizational design. You can have the best models and infrastructure in the world, but without the right team structure, operating model, and governance, AI initiatives stall in pilot purgatory, fight for resources, or die when a single "AI hero" leaves.

The data is unambiguous. Organizations with formal AI Centers of Excellence are up to 3× more likely to move from pilots to scaled AI programs. Companies that treat AI as a centralized capability with shared standards report 35% lower technology expenses and 60% higher impact than those running AI BU-by-BU. The Chief AI Officer role has nearly tripled in adoption as boards seek a single accountable executive for AI strategy, ethics, and ROI.

This playbook is a complete, leadership-grade operating manual: how to choose your team structure, what an AI Center of Excellence should actually do, which roles you really need (and in what order), how AI product management fits in, what to measure, and a 12-month roadmap to build it all. Use it for board-level planning, quarterly business reviews, hiring plans, and as a basis for your own internal charter.

> **Three structural commitments for 2026**
>
> **1. A named executive owner for AI** — CAIO, CDO, or equivalent — with real P&L exposure and decision rights.
> **2. A Center of Excellence** that owns shared platforms, governance, and reusable assets — not a science lab.
> **3. Embedded AI product managers** with KPI accountability tied to business outcomes, not feature ship dates.

How to use this playbook

Read Parts 1–4 before your next leadership offsite to align on operating model and structure. Use Part 5 (12-month roadmap) and Part 6 (KPI dashboard) as the spine of your AI program plan. Run the Part 7 checklist quarterly with the AI Steering Committee. Adapt the templates in Part 8 directly into your governance tool of record.

Contents

Part 1 — The AI Operating Model

Thereis noone "correct" AI orgchart. The right structure depends on your maturity, scale, use-case mix, regulatory burden, and how harmonized your data and tech stacks are across business units. The single most common mistake is starting too federated — fragmenting data, duplicating platforms, and stalling governance.

1.1 Three operating models

Model	How It Works	Best Fit	Risk if Mis-applied
Centralized CoE	A single AI CoE owns strategy, platforms, best practices, and most delivery. BUs are stakeholders and consumers.	Early-stage AI maturity, regulated industries, when governance is weak	Bottleneck; CoE becomes order-taker; BUs feel disempowered
Federated / Hub & Spoke	Central CoE sets standards, platforms, and governance. Embedded AI pods in BUs handle domain-specific delivery.	Diversified enterprises with multiple product lines, regions, or verticals	Drift between BUs; standards diverge; governance becomes advisory only
Embedded / Decentralized	Data and AI teams sit entirely within BUs. Minimal central coordination; governance handled ad hoc.	Very mature orgs with strong, aligned BU leaders — usually combined with a light central layer	Duplicated effort, inconsistent risk posture, audit failures

1.2 Decision criteria

Use this matrix to choose your starting structure:

Criterion	Centralize	Federate	Embed
High-value AI use cases in scope	< 5	5–20	20+
Data & tech stack consistency across BUs	Fragmented	Mixed	Harmonized
Existing data governance maturity	Weak	Developing	Strong
Regulatory / compliance burden	High	Moderate	Low
BU leader AI fluency	Low	Mixed	High
Reuse potential across BUs	High	Medium	Low
Tolerance for delivery latency	Low	Medium	High

1.3 The evolution pattern (Year 1 → Year 3)

Horizon	Structure	CoE Role	BU Role
Year 1 — Build	Centralized CoE	Strategy, platforms, governance, most delivery	Sponsor lighthouse use cases; provide SMEs
Year 2 — Spread	Federated / Hub & Spoke	Platforms, patterns, governance, complex use cases	Embedded pods deliver domain use cases

Horizon	Structure	CoE Role	BU Role
Year 3 — Scale	Federated / Light embed	Platforms, governance, AI strategy, enablement	Most delivery; CoE supports specialised work

The single biggest operating-model mistake

Trying to start federated. Without shared platforms, patterns, and governance in place first, federated structures fragment data, replicate infrastructure, and create governance gaps that surface in audit cycles. **Centralize early. Federate as capabilities mature.**

1.4 Common operating-model pitfalls

Pitfall	Symptom	Fix
CoE as a science lab	Impressive prototypes, few production wins, BUs disengage	Tie CoE charter to business outcomes; embed AI PMs; require business sponsors
Wildcard delegation	Each BU runs its own AI; no shared platform; conflicting policies	Stand up minimum-viable CoE: platform, governance, intake, standards
Governance bolted on late	Audit failures, model risk surprises, regulator engagement	Set up Risk Council in Phase 1; gate production deploys on review
Hero-driven delivery	One brilliant person owns the AI program; bus factor of one	Document patterns, rotate ownership, fund a 2-deep model on every system
Vague accountability	Multiple execs claim AI; nobody is accountable for ROI	Name a single executive owner with P&L exposure

1.5 Industry variations

Industry		Critical Add-ons
Financial services	Centralized CoE → Federated by Year 2	Model risk function (SR 11-7), validators independent of builders, evidence packs by default
Healthcare / life sciences	Centralized CoE	Clinical safety officer, IRB liaison, dataset provenance, regulator-grade documentation (e.g. FDA SaMD)
Retail / consumer	Federated by category early	Personalization platform team; experimentation infra; consent & preference management
Manufacturing / industrial	Centralized CoE with site pods	OT/IT integration; edge deployment; safety-critical review boards
Public sector	Centralized CoE with strict governance	Algorithmic transparency register; public-impact assessment; procurement & vendor lock-in controls

Part 2 — AI Center of Excellence Charter

An AI CoE is more than "a teamof smart datascientists." Itisa cross-functionalunitthat sets AI strategy, owns shared platforms, enforces governance, and accelerates delivery across the enterprise. Use the structure below as a charter template you can adapt and sign.

2.1 Mission & scope

CoE Mission Statement (template)

To make AI a repeatable, governed, value-creating capability across the enterprise — by setting strategy, providing shared platforms and patterns, enforcing responsible-AI standards, and accelerating delivery of high-impact use cases in partnership with business units.

Scope (in)	Scope (out)
AI strategy and portfolio management	BU-specific product roadmap (BUs own)
Shared data & AI platforms (feature store, registry, RAG infra)	BU-specific operational systems
AI governance, model risk, ethics policy	Enterprise-wide IT security policy (CISO owns)
Lighthouse delivery and complex use cases	Routine analytics and BI (BUs own)
AI literacy, communities of practice	General training programs (HR owns)
Vendor/partner evaluation for AI tooling	General procurement (Procurement owns)

2.2 Five core functions

Strategy & Portfolio Management

Map AI opportunities to business goals; maintain an AI portfolio (pipeline, in-flight, in production); define value metrics and track ROI.

- Annual AI strategy aligned with corporate plan
- Quarterly portfolio review with steering committee
- Prioritized backlog with business sponsors
- Investment cases for top initiatives

Platform & Architecture

Provide shared data and AI platforms; define reference architectures, guardrails, and the path from notebook to production.

- Shared feature store, model registry, experiment tracking
- RAG / retrieval infrastructure with consent and residency tags
- Reference architectures for top patterns (RAG, copilot, classification, agent)
- Self-serve developer experience with guardrails baked in

Delivery & Enablement

Build and deploy high-impact use cases; provide reusable components; train and support BU teams and citizen developers.

- 3–5 lighthouse use cases in production by month 9
- Reusable components, templates, and accelerators
- BU enablement playbooks and SME pairing
- Citizen-developer guardrails and approved tools

Governance & Risk

Implement AI governance frameworks (ISO 42001, NIST AI RMF, EU AI Act); run model risk assessments, validation, and monitoring; enforce policies for data, security, ethics, and compliance.

- Model risk policy and tiering scheme
- Pre-deployment review and evidence packs
- Production model inventory + MBOM
- Audit-ready evidence on demand (< 5 days)

Talent & Culture

Hiring, upskilling, and AI literacy; foster communities of practice; build the learning culture that AI work requires.

- AI literacy program for executives, PMs, engineers
- AI Practitioners Guild — bi-weekly knowledge sharing
- Hiring partnerships and talent pipeline plan
- Career ladders for ML/AI/data roles

2.3 Governance bodies

Body	Cadence	Purpose	Members
AI Steering Committee	Monthly	Strategy, portfolio, escalations, KPI review	CAIO (chair), CIO/CTO, CDO, CFO, CRO, BU leaders
AI Architecture Board	Bi-weekly	Reference architectures, platform decisions, exceptions	Head of AI Platform, ML/AI engineers, security architect
AI Risk Council	Monthly	Model risk, ethics review, compliance posture	Governance lead, legal, privacy, CISO, CAIO
AI Practitioners Guild	Bi-weekly	Knowledge sharing, patterns, code review, retros	ML engineers, data scientists, AI PMs, data engineers
Use-Case Intake Forum	Weekly	Triage new AI requests; assign owners; risk-tier	AI PMs, governance lead, BU sponsors

Sample AI Steering Committee agenda (60 min)

Time	Topic	Owner
0–5	Confirm agenda & previous actions	Chair (CAIO)
5–15	Portfolio & KPI dashboard review	CoE Lead
15–30	Use-case approvals & major risks	AI PMs / Risk Lead
30–45	Investment & funding decisions	CFO / CAIO
45–55	Strategic update or deep-dive	Rotating
55–60	Decisions log & next steps	Chair

2.4 Funding models

Model	How It Works	Pros	Cons	Use When
Central	CoE funded by corporate / IT budget; BUs use services free at point of use	Predictable; encourages adoption; clear accountability	Risk of being seen as cost centre; no demand signal	Year 1 build phase; standardisation critical
Charge-back	BUs pay for services rendered (delivery, platform, governance reviews)	Strong demand signal; aligns priorities	Heavy admin; may discourage early adoption; pricing fights	Mature, federated programs
Hybrid	Platform & governance centrally funded; delivery charged-back to BUs	Best of both — adoption at platform layer, demand-driven delivery	Allocation rules need clarity	Year 2+; most enterprises end here

2.5 Decision rights

Makedecisionrightsexplicit. Below is a starting set — adapt to your matrix structure and document in your charter:

Decision	CAIO	CoE	BU	Risk
AI strategy & portfolio priorities	Decide	Recommend	Input	Input
Platform standards & reference architectures	Approve	Decide	Input	Input
Use-case approval (low risk)	Inform	Decide	Sponsor	Inform
Use-case approval (high risk)	Decide	Recommend	Sponsor	Approve
Vendor / tooling selection (AI)	Approve	Decide	Input	Input
Production go-live	Inform	Decide	Sponsor	Approve
Model retirement / kill	Inform	Decide	Input	Recommend

Part 3 — The AI Role Library

Hiring"afewdatascientists"is not enough.SuccessfulAI orgs combine technical, product, domain, and governance roles. Use this library as a hiring spec and to identify gaps in your current team.

3.1 Executive layer

Role	Owns	Reports To	Time on AI
Chief AI Officer (CAIO)	AI strategy, portfolio, ethics, ROI; chairs Steering Committee	CEO or COO	100%
CIO / CTO	Infrastructure, platforms, integration, engineering standards	CEO	30–50%
CDO	Data quality, governance, lineage, master data, data product strategy	CEO or COO	40–60%
CFO	AI investment decisions, ROI tracking, business case approval	CEO	10–20%
Chief Risk / Legal / Ethics	AI risk policy, compliance, responsible-AI guardrails	CEO or Board	20–40%
CHRO	AI talent strategy, literacy programs, change management	CEO	20–30%

3.2 Core delivery roles

Role	Mission	Core skills	KPI focus
Data Engineer	Builds and maintains data pipelines and storage. Ensures data is reliable, timely, and ready for modeling and retrieval. Owns ingestion, transformation, lineage, and freshness SLAs.	SQL, Python, Spark/dbt, Airflow, lakehouse/warehouse, schema design	Pipeline reliability, data freshness, lineage coverage
Data Scientist / ML Researcher	Explores data, frames problems, designs experiments, builds models. Works closely with domain experts. Owns offline evaluation.	Statistics, ML, Python, experiment design, feature engineering	Model accuracy on business metric, experiment throughput
ML Engineer / AI Engineer	Productionizes models; handles serving, scaling, monitoring; integrates AI into apps. Owns the path from notebook to production.	MLOps, serving frameworks, observability, software engineering	Time-to-prod, production reliability, error budgets
MLOps / Platform Engineer	Builds the shared ML platform: CI/CD for models, feature stores, model registries, observability. Standardizes dev → prod.	Cloud infra, IaC, Kubernetes, MLflow/Vertex/SageMaker, monitoring	Platform adoption, reuse rate, deployment speed
AI Product Manager	Owns problem definition, user experience, and business outcomes. Balances feasibility, value, ethics, and risk. Designs evaluation frameworks (offline and online).	Product thinking, eval design, prompt design, RAG vs fine-tune trade-offs	Use-case ROI, adoption, business KPI movement

Role	Mission	Core skills	KPI focus
Domain Expert / UX Designer	Provides real-world constraints; ensures solutions fit workflows. Designs human-AI interactions, confidence cues, recovery paths.	Domain knowledge, interaction design, prototyping, user research	Adoption, task success rate, user trust
AI Governance / Risk Lead	Embeds responsible-AI, compliance, and risk controls into the lifecycle. Runs model risk reviews; maintains the model inventory and MBOM.	Risk management, regulatory frameworks (ISO 42001, NIST, EU AI Act)	Risk reviews completed, audit readiness, incidents avoided

3.3 Stage-appropriate hiring

Stage	Use Cases	Core Hires (in order)	Partner For
Early	0–3	Head of AI + Data Engineer → ML Generalist (shared) → AI PM	MLOps, design, advanced modeling, governance
Growth	3–10	Platform/MLOps Eng → ML Engineers (2–3) → Governance Lead → Data Scientists → AI PMs (2)	Specialist domain modeling, niche infra, surge capacity
Scale	10+	Heads of (Platform / Delivery / Governance) → embedded BU pods (ML eng + data eng + PM × N) → security & risk specialists	Vertical expertise, surge capacity, regulator interaction

3.4 Build vs Buy vs Borrow

Capability	Default	Why
AI strategy & portfolio	**Build**	Strategic; cannot be outsourced
AI product management	**Build**	Owns the link between business and AI
Data engineering for AI	**Build**	Tightly coupled to your data estate
ML / AI engineering (core)	**Build**	Production reliability is a learned skill
Specialised modeling (NLP, vision, forecasting)	Buy / Borrow	Spiky demand; vendor or partner often faster
MLOps platform	Buy with build wrappers	Vendor platforms cover 80%; you customise
Governance & risk reviews	Build (with advisory)	Regulator-facing; cannot delegate
AI literacy training	Borrow / Buy	Mature off-the-shelf programs exist

3.5 Indicative compensation bands

Bands below are illustrative forguidanceonly —calibrate to your region, industry, and benchmarking source. All figures USD, base + target bonus, mid-career level (5–10 years' experience):

Role	US (Hub Cities)	UK / EU	India
Head of AI / CoE Lead	$280–450k	£200–320k	70–1.4Cr
AI Product Manager	$200–320k	£140–220k	■40–80L
Senior ML Engineer	$220–350k	£150–230k	■45–90L
MLOps / Platform Engineer	$200–320k	£140–210k	■40–80L
Data Engineer	$170–260k	£120–180k	■30–60L
Data Scientist	$170–270k	£110–180k	■30–60L
AI Governance Lead	$180–280k	£130–200k	■35–70L

The most common hiring mistake

> Over-hiring researchers and under-hiring engineers and PMs. Result: impressive prototypes, few production wins. For every data scientist, plan for at least one ML engineer, one data engineer, and fractional product and governance ownership.

Part 4 — Org Charts by Stage

Thesereferenceorg chartsshow how the structureevolves asthe AI portfolio grows. Use them as a starting point — adapt headcount and reporting lines to your industry and context.

4.1 Early stage (0–3 use cases)

Lean, founder-mode AI team. Standardise tools and patterns from day one to avoid rework when you scale. Lean on partners for niche skills.

Layer	Roles	Headcount
Sponsor	Executive sponsor (CDO, CTO, or CEO)	0.2 FTE
Lead	Head of AI / AI lead	1
Build	Data/ML generalist, Data engineer, AI PM (shared)	2–3
Govern	Privacy/legal partner (matrixed)	0.2 FTE
Total		**3–4 FTE**

Goal: ship 1–2 lighthouse use cases end-to-end and prove governance can scale.

4.2 Growth stage (3–10 use cases)

Stand up the CoE. Hire the platform engineer before the next data scientist. Begin embedding pods in 1–2 BUs.

Layer	Roles	Headcount
Sponsor	AI Steering Committee (CAIO/CTO chair)	—
CoE Lead	Head of AI / CoE Lead	1
Platform	MLOps / Platform engineer, Data platform engineer	2
Delivery	ML engineers, Data engineers, Data scientists	5–7
Product	AI Product Manager(s)	2
Govern	Governance / Risk lead, partnered with legal & CISO	1
Total		**11–13 FTE**

Goal: stand up CoE and shared platform; ship 3–5 use cases in production; embed pods in 1–2 BUs.

4.3 Scale stage (10+ use cases)

AI as a factory. Central CoE owns platforms, standards, governance; BU pods own domain delivery; specialists in risk, security, ethics.

Layer	Roles	Headcount
Sponsor	CAIO + AI Steering Committee	—
CoE Core	CoE Lead, Head of Platform, Head of Delivery, Head of Governance	4

Layer	Roles	Headcount
Platform	MLOps, Data platform, Security, Observability	5–7
Delivery	Senior ML/AI engineers, data engineers, data scientists	10–15
Product	Centralized AI PMs + embedded BU PMs	4–6
Embedded	BU pods (ML eng + data eng + PM, ×N)	8–20
Govern	Governance, Risk, Ethics, Privacy partners	3–4
Total		**35–55+ FTE**

Goal: AI as a repeatable, governed, instrumented factory across the enterprise.

4.4 Industry-specific patterns

Theseaugmentationslayerontopofthestage-based structures above:

Industry	Augmentations	Critical roles to add
Financial services	Independent model validation; SR 11-7 alignment; evidence packs by default; concentration limits on third-party models	Model Risk Manager, Independent Validators, Regulator Liaison
Healthcare / life sciences	Clinical safety review board; IRB liaison; FDA SaMD pathway where applicable; clinician-in-the-loop by design	Clinical Safety Officer, Bioethics Lead, Medical Affairs partner
Retail / consumer	Personalization platform; experimentation infra; real-time ML; consent and preference management	Personalization Lead, ExperimentOps, Customer Data Steward
Manufacturing / industrial	OT/IT integration; edge deployment; safety-critical review; downtime & quality KPIs	OT Integrator, Edge ML Engineer, Safety Engineer
Public sector	Algorithmic transparency register; public-impact assessments; procurement & vendor lock-in controls	Algorithmic Transparency Lead, Public Engagement, Procurement Specialist

Part 5 — 12-Month Roadmap

The roadmapbelow assumes youare startingfromscattered AIexperiments and want to reach a functioning CoE plus 3–5 production use cases inside 12 months. Each phase has clear deliverables, RACI, and exit criteria.

Phase 1 — Months 0–3: Assess & Design

Objective: understand current state, align leadership, and design the target AI organization.

Workstream	Activities	Owner
Inventory	Map current AI initiatives, tools, people; capture spend and risks	Head of AI (interim)
Gap analysis	Identify gaps: data engineering, MLOps, product, governance	Head of AI
Operating model	Decide centralized vs hybrid; draft CoE charter; define RACI	CAIO + Steering Cmte.
Governance	Stand up AI Steering Committee; define escalation paths	CAIO + CRO
Use-case shortlist	Identify 5–8 candidates; risk-tier; prioritise 2–3 for Phase 2	AI PM + BU sponsors
Talent plan	Define Phase-2 hiring plan; identify partners	CAIO + HR
Deliverables	**AI org blueprint, RACI, CoE charter v0.9, prioritized use-case shortlist, talent plan, governance charter**	—

Weekly milestones — Phase 1

Wk	Milestone
1–2	Kickoff; appoint Head of AI (or interim); align sponsors
3–4	Inventory complete; first stakeholder interviews
5–6	Gap analysis; operating-model options memo
7–8	CoE charter v0.5; AI Steering Committee stood up
9–10	Use-case shortlist agreed; risk-tier complete
11–12	Charter v0.9 signed; Phase-2 plan and budget approved

Phase 2 — Months 3–6: Stand Up the CoE & Core Platform

Objective: stand up a functional CoE and shared platform; begin lighthouse delivery.

Workstream	Activities	Owner
Hiring	Hire core CoE: 1–2 senior ML/AI engineers, 1 data engineer, 1 AI PM, 1 governance lead, 1 platform engineer	CAIO + HR
Platform	Stand up data & ML platform: MLOps, experiment tracking, model registry, basic RAG infra	Head of Platform
Standards	Reference architectures for top patterns; security & compliance baseline; intake form	Architecture Board
Lighthouse projects	Kick off 2–3 lighthouse use cases with named business sponsors	AI PMs + BUs
Governance	Risk policy live; pre-deployment review process operational	Risk Lead
Deliverables	**Operational CoE; first shared services; 2–3 lighthouse use cases in build; risk & intake processes live**	—

Phase 3 — Months 6–9: Embed & Federate

Objective: extend AI capability into BUs without losing coherence.

Workstream	Activities	Owner
BU pods	Create embedded AI pods in 1–2 BUs (e.g., customer ops, risk)	BU leaders + CAIO
Standards	CoE provides playbooks, reference architectures, and platform; pods own domain delivery	CoE
Governance	Formalize model risk reviews, ethics checks, architecture exception process	Governance lead
Reuse	Identify and harvest reusable components; document in pattern catalog	Architecture Board
Adoption	Run AI literacy program for execs and BU leaders	CAIO + CHRO
Deliverables	**3–5 use cases live across multiple BUs; pattern catalog; literacy program; first BU pods operational**	—

Phase 4 — Months 9–12: Industrialize & Optimize

Objective:treat AI asafactory —repeatable, governed, and measurable.

Workstream	Activities	Owner
KPIs	Implement CoE KPI dashboard; quarterly business review	CoE Lead
Portfolio	Introduce AI portfolio management — prioritization, resource allocation	CAIO
Talent	Strengthen pipelines, communities of practice, and AI literacy training	CHRO + CoE
Optimization	Cost optimization; consolidate tools; retire models that no longer pay back	Head of Platform
Federate	Plan and stage transition to more federated model where appropriate	CAIO
Deliverables	**AI "factory" cadence; transparent KPI dashboard; transition plan to federated model; year-2 plan and budget**	—

5.5 Risk register

Tracktheseasastanding item in your monthly Steering Committee:

Risk	Likelihood	Impact	Mitigation
Hero dependency: program collapses if key person leaves	High	High	Two-deep on every system; pair-programming; documented patterns
Pilot purgatory: many pilots, few in production	High	High	Force production gate at 90 days; require business sponsor
Governance bolted on late: audit / regulator surprise	Medium	High	Risk Council from Phase 1; pre-deploy review mandatory
Vendor lock-in & cost spike on tokens / APIs	Medium	Medium	Multi-provider abstraction; cost dashboards; rate limits
Data residency / cross-border breach	Medium	High	Region-pinned platforms; egress allow-lists; residency tags
Talent attrition (key role)	Medium	High	Retention plan; equity refreshers; clear career ladders
BU push-back ("the CoE is slowing us down")	Medium	Medium	Clear SLAs; embedded PMs; opt-in self-serve paths

Part 6 — Executive KPI Dashboard

Ifyour CoEcannotshow impact onatleastthreeoftheseKPIs, it's atriskof being seen as a science lab rather than a strategic function. Track monthly; review quarterly with the AI Steering Committee.

KPI	Definition	2026 Target	Benchmark
Strategy			
AI initiatives aligned to business goals	% of in-flight projects mapped to a corporate goal	≥ 90%	Top quartile: 95%+
Use cases in production	Cumulative count of AI use cases live in production	≥ 5 by month 12	Top quartile: 8+
Pipeline health	Use cases at each lifecycle stage (idea, pilot, prod, retired)	Balanced (no > 50% stuck)	—
Delivery			
Deployment speed	Days from approved idea to production	< 90 days	Top quartile: < 60
Reuse rate	% of new projects using ≥ 1 shared asset	≥ 60%	Top quartile: 75%+
Production reliability	Incidents per model per quarter	< 1	Top quartile: ~0.3
Mean time to recover (MTTR)	Average minutes to recover from a model incident	< 60 min	—
Business impact			
AI-attributed ROI	Annualized run-rate from AI use cases	≥ 2× CoE cost	Top quartile: 4×+
Adoption	Active users / target population per use case	≥ 60%	Top quartile: 80%+
Time saved per user	Hours/week saved attributable to AI	Track per use case	—
Governance & risk			
Models with completed risk review	% in production with completed pre-deploy review	100%	100%
High-risk findings open > 30 days	Count	0	0
Audit readiness	Time to produce model evidence pack on request	< 5 days	—
AI incident severity distribution	Count of S1/S2/S3 incidents per quarter	0 S1	—
Talent & culture			
AI literacy coverage	% of execs & PMs with completed training	≥ 80%	—

KPI	Definition	2026 Target	Benchmark
AI practitioner attrition	Annualized voluntary attrition	< 10%	Industry avg: ~12%
Internal mobility into AI roles	% of AI hires sourced internally	≥ 25%	—

Part 7 — 30-Point AI Org Readiness Checklist

Run thischecklist quarterly withthe AI SteeringCommittee."Critical" items must bein place before scaling; "Important" should be on the 12-month plan. If you can tick most boxes — or see a clear path in 12 months — your organization is structurally ready to scale AI.

Strategy & Leadership

#	Check	Priority	Owner
1	Named executive owner for AI (CAIO or equivalent) with P&L exposure	Critical	_____
2	■ AI strategy aligned with business goals and approved by leadership	Critical	_____
3	■ AI Steering Committee meeting at least monthly with KPI review	Critical	_____
4	■ Clear view of top 10 AI initiatives, sponsors, and value	Important	_____
5	■ Annual AI investment plan tied to business outcomes	Important	_____

Structure & CoE

#	Check	Priority	Owner
1	Decision made on centralized vs federated vs hybrid model	Critical	_____
2	■ AI CoE charter documented and signed	Critical	_____
3	■ CoE staffed across tech, product, governance	Critical	_____
4	■ At least one BU has an embedded AI pod	Important	_____
5	■ Clear interfaces (RACI, intake, escalation) between CoE and BUs	Important	_____

Roles & Hiring

#		Priority	Owner
1	Dedicated data engineering capacity for AI projects	Critical	_____
2	■ ML / AI engineers responsible for productionization	Critical	_____
3	■ AI product managers with KPI ownership	Critical	_____
4	■ Governance / risk lead embedded from design through deployment	Critical	_____
5	■ Stage-appropriate mix of in-house and partner talent	Important	_____

Processes & Platforms

#	Check	Priority	Owner
1	■ Standard AI lifecycle from idea → pilot → production → monitoring	Critical	_____
2	■ Shared data / ML platform (experiment tracking, registry, RAG infra)	Critical	_____
3	■ Pre-deployment risk review mandatory and documented	Critical	_____
4	■ Playbooks for common patterns (RAG, copilot, classification, agent)	Important	_____

#	Check	Priority	Owner
5	Audit-ready evidence packs available on demand (< 5 days)	Important	_____

Data & Governance

#	Check	Priority	Owner
1	AI model inventory + Model Bill of Materials (MBOM) maintained	Critical	_____
2	■ Data residency & classification tags applied to AI workflows	Critical	_____
3	■ Consent / preference management integrated with AI pipelines	Important	_____
4	■ Erasure / right-to-erasure pipelines for AI stores	Important	_____
5	■ Vendor / third-party AI risk reviewed at procurement	Important	_____

Culture & Enablement

#	Check	Priority	Owner
1	AI literacy programs for executives and key functions	Critical	_____
2	■ Communities of practice / guilds for AI practitioners	Important	_____
3	■ Clear guidelines on acceptable AI use and risk thresholds	Critical	_____
4	■ Recognition and incentives linked to successful AI adoption	Important	_____
5	■ Knowledge base of patterns, mistakes, and reusable assets	Important	_____

Part 8 — Templates & Tools

Copythesetemplatesintoyourgovernancetoolofrecord.Adapt the wording to your charter; keep the structure.

A. CoE Charter (one-pager)

Field	Content
Mission	One sentence — why this CoE exists.
Scope (in)	What the CoE owns end-to-end.
Scope (out)	What stays with BUs / other functions.
Sponsor	Executive accountable for outcomes.
Lead	Day-to-day owner.
Operating model	Centralized / Federated / Hybrid.
Year-1 priorities	3–5 measurable goals.
Year-1 KPIs	Strategy, delivery, governance metrics with targets.
Funding model	Central, charge-back, or hybrid; allocation rules.
Decision rights	What the CoE decides vs. recommends vs. supports.
Governance bodies	Steering / Architecture / Risk / Guild — cadence and membership.
Review cadence	Monthly KPIs; quarterly portfolio; annual charter refresh.

B. AI Use-Case Intake Form

Field	Description
Use-case name	Short descriptive title.
Business sponsor	Named exec accountable for outcome.
Problem statement	What changes for the business if we ship this?
Target user / workflow	Who interacts with the AI; where in the process.
Success metrics	Leading + lagging; quantitative.
Data needed	Sources, sensitivity, residency, retention.
AI pattern	RAG, classification, agent, copilot, forecasting, etc.
Risk classification	Low / Medium / High (per AI risk policy).
Governance gates required	Risk review, ethics review, legal sign-off.
Estimated effort	T-shirt size (S/M/L/XL) and FTE-weeks.
Build vs Buy	If buy: shortlist of vendors with rationale.
Decision	Approve / Defer / Reject — with rationale.

C. AI RACI (sample)

Activity	CAIO	CoE	BU	Risk	Legal
AI strategy	A	R	C	C	C
Use-case prioritization	A	R	C	I	I
Platform & standards	A	R	C	C	I
Use-case build & deploy	C	S	R	C	C
Model risk review	A	S	C	R	C
Audit & evidence	A	S	S	R	C
Adoption & training	A	S	R	I	I
Vendor / partner selection	A	R	C	C	C
Incident response	A	R	C	C	I

R = Responsible · A = Accountable · C = Consulted · I = Informed · S = Supports

D. Role JD Template

Field	Content
Title & level	e.g., Senior ML Engineer, IC4
Reports to	Head of AI Platform
Mission	Why this role exists in one sentence.
Outcomes (Year 1)	3–5 measurable outcomes the role is hired to deliver.
Responsibilities	5–8 bulleted areas of work.
Must-have skills	Technical and behavioural — required.
Nice-to-have	Differentiators — preferred.
Interview loop	Stages, owners, calibration approach.
KPIs	How we measure success in this role.
Compensation band	Base + bonus + equity; level mapping.

E. AI Steering Committee Agenda (60 min)

Time	Topic	Inputs	Decisions / Outputs
0–5	Confirm agenda & previous actions	Action log	Updated log
5–15	KPI dashboard review	Monthly KPI pack	Highlight top 3 issues
15–30	Use-case approvals & risks	Intake forum recommendations, risk register	Approve / defer / reject; risk acceptances
30–45	Investment & funding decisions	Business cases, budget tracker	Approve / defer / reject; commit budget
45–55	Strategic update or deep-dive	Topic memo (rotating)	Direction agreed
55–60	Decisions log & next steps	—	Owners, dates, escalations

F. AI Investment Business Case (template)

Section	Content
Problem & opportunity	What changes for the business; quantified.
Proposed solution	Use case, AI pattern, scope of MVP.
Value (3-year)	Revenue uplift / cost saved / risk reduced; baseline + assumptions
Investment	FTE, infra, vendor, change management — by year
Payback & NPV	Discount rate, payback period, NPV
Risks & mitigations	Top 5 risks, owners, controls
KPIs & milestones	Leading + lagging; review cadence
Decision sought	Approve / Defer / Reject — and budget envelope

G. Vendor / Partner Evaluation Scorecard

Criterion	Weight	Score 1–5	Notes
Capability fit (must-haves)	25%	_	
Architecture & integration fit	15%	_	
Data residency & security posture	15%	_	
Governance & auditability	10%	_	
Total cost of ownership (3-year)	15%	_	
Implementation & support model	10%	_	
Reference customers & track record	5%	_	
Lock-in & portability risk	5%	_	
Total	100%	_	

H. Quarterly Business Review (QBR) template

Section	Content
1. Headline (1 slide)	3 KPIs that matter, 1 win, 1 risk, 1 ask of leadership
2. Strategy & portfolio	Use-case lifecycle funnel; what shipped, what's in build, what's blocked
3. KPI dashboard	Strategy, delivery, business impact, governance, talent — vs. target
4. Wins (case studies)	Top 2–3 use cases this quarter — value realised, lessons learned
5. Risks & incidents	Risk register status; any S1/S2 incidents; root cause & CAPA
6. Investment & spend	Burn vs. plan; tooling cost; vendor concentration
7. Talent	Hiring pipeline; attrition; literacy program coverage
8. Outlook	Next quarter's plan: top 5 deliverables, risks, decisions needed
9. Decisions & asks	Specific decisions or unblocking required from steering committee

Part 9 — Industry Patterns

Theseshort patternscapturethestructuraldifferencesAIorganizations in each industry tend to need on top of the generic playbook.

Financial services

Regulators expect explainability, validation independence, and evidence-on-demand. Independent model risk function is mandatory; concentration limits on third-party models; SR 11-7-aligned documentation by default.

- Independent Model Risk Manager (separate from builders)
- Pre-prod and quarterly validation of high-risk models
- Concentration limits and exit plans for third-party providers
- Regulator liaison and evidence-pack-by-default workflow

Healthcare / life sciences

Patient safety and regulatory pathways (e.g., FDA SaMD) reshape every stage. Clinician-in-the-loop is the default for any patient-impacting decision; clinical safety officer signs off.

- Clinical Safety Officer with veto on patient-impacting AI
- IRB liaison for any human-subject AI
- Dataset provenance and consent at the chart level
- Versioned, frozen datasets for any clinical-grade model

Retail / consumer

Personalization is the centre of gravity. Real-time ML, experimentation, and consent management dominate. CDP integration is the spine of the AI platform.

- Personalization platform team (often the largest pod)
- ExperimentOps function for A/B and bandit infrastructure
- Consent & preference engine wired into every model call
- Real-time feature engineering at sub-second latency

Manufacturing / industrial

OT/IT integration and edge deployment dominate. Models often run on constrained edge devices with safety-critical implications.

- OT integrator with shop-floor and PLC expertise
- Edge ML engineer for constrained devices
- Safety review board for any control-loop AI
- Downtime/quality KPIs as primary value metrics

Public sector

Publicaccountability, procurement constraints, and risk-aversion to vendor lock-in shape the operating model. Algorithmic transparency is non-negotiable.

- Algorithmic transparency register (public-facing)
- Public-impact assessments before high-risk deployments
- Open-standards-first procurement strategy
- Plain-language explanation requirements for all AI decisions

Part 10 — Anti-Patterns & Pitfalls

Theseare thestructuralfailures we see most often inenterpriseAI programs. Each one has predictable symptoms — recognise them early and act before they become entrenched.

Anti-Pattern	Symptom	Fix
The science lab	Impressive prototypes; few production wins; BUs disengage	Tie CoE charter to business outcomes; embed AI PMs; require sponsors
Hero-driven program	One brilliant person owns everything; bus factor of one	Document patterns; pair-program; fund a 2-deep model on every system
Wildcard delegation	Each BU runs its own AI; conflicting policies; duplicated infra	Stand up minimum-viable CoE: platform, governance, intake, standards
Vague ownership	Multiple execs claim AI; nobody is accountable for ROI	Name a single executive owner with P&L exposure
Governance bolted on late	Audit failures, model risk surprises, regulator engagement	Stand up Risk Council in Phase 1; gate production deploys on review
Tool-first thinking	We bought the platform; now what? No use cases, no value	Use case → architecture → tools — never the reverse
Pilot factory	20 pilots, 2 in production; nothing scales	Force production gate at 90 days; require sponsor and exit criteria
Centralization forever	CoE becomes a bottleneck; BUs hate the queue	Federate when capabilities mature; build self-serve paths early
Boundary-less CoE	CoE owns everything; nothing is a BU's job; no scaling	Tight scope-out; clear interfaces (RACI, intake, escalation)
Talent monoculture	All researchers, no engineers; or all engineers, no domain experts	Balance: every team has product, engineering, domain, and governance

Part 11 — Frequently Asked Questions

Q:ShouldwestartwithacentralizedCoEorembedAIteamsdirectlyineachbusiness unit?

A: Most organizations benefit from a centralized CoE at first, especially when data, tooling, and governance are fragmented. As maturity grows, embedded teams in BUs can take on more responsibility, while the CoE focuses on platforms, standards, and governance.

Q: Do we really need a Chief AI Officer?

A: Not every company needs a distinct CAIO, but you do need a clearly accountable executive for AI outcomes. In some organizations this is the CDO or CTO; in others, a dedicated CAIO simplifies accountability and signals strategic commitment to the board and the market.

Q: How big should our AI CoE be?

A: Early CoEs often start with 5–10 people, then grow as demand and value increase. Focus less on headcount and more on coverage of key functions: strategy, data/platform, delivery, governance, and enablement.

Q: What work should stay in-house vs. be outsourced?

A: Keep strategy, data governance, and core product ownership in-house. Consider partnering for specialised skills (e.g., advanced MLOps, domain-specific modeling) or when you need to move quickly without long-term headcount commitments.

Q: How do we prevent AI teams from becoming a "science lab" disconnected from the business?

A: Give AI teams clear KPIs tied to business outcomes, embed AI product managers, and require business sponsors for each initiative. Make ROI and adoption part of performance reviews, not just technical metrics.

Q: How do we handle regulators (GDPR, EU AI Act, sector regs)?

A: Stand up the Risk Council in Phase 1. Require evidence packs by default. Engage your DPO, CISO, and legal team in the operating-model design — not after the first pilot ships. For high-risk use cases, pre-engage regulators where the regulatory framework permits.

Q: How do we measure AI ROI?

A: Use a value-tree per use case: business KPI movement → adoption rate → model output quality → infrastructure cost. Track leading and lagging metrics. Aggregate across the portfolio for the CoE-level ROI ratio (target: ≥ 2× CoE annual cost).

Q: How should we think about Build vs Buy for AI platforms?

A: Buy the foundation, build the differentiation. Vendor platforms cover ~80% of MLOps needs; build custom wrappers for your data estate, patterns, and governance. Avoid building generic capabilities; avoid buying your strategic moats.

Make This Real — Run an AI Org & CoE Design Workshop

Theframeworks inthisplaybooktakemonthsof debateto applyfroma blank slate. A focusedworkshop with your leadership team compresses those months into a few days of decisions you can act on the same week.

Day	Output
Day 1 — Diagnose	Capability assessment across structure, roles, processes, governance, platform, talent. Output: **Capability map** with gaps and quick wins.
Day 2 — Design	Operating model decision; CoE charter v1; org chart by stage; year-1 KPIs. Output: **Signed CoE charter** and **Year-1 plan**.
Day 3 — Decide	Hiring plan, partner shortlist, investment ask, 12-month roadmap with quarterly milestones. Output: **Approved roadmap and budget envelope**.

Book the workshop

Two to three days. Your leadership team. Our experts. You leave with a signed CoE charter, a year-1 roadmap, an approved hiring plan, and investment envelope you can act on the same week.

Contact: hello@fracto.ie · **Web:** fracto.ie · **Series:** Enterprise AI Playbooks 2026