FRACTO

AI in Financial Services
Blueprint

Use-Case Maps·Reference Architectures · Model Governance · EU AI Act

WHAT'S INSIDE

- Banking & insurance AI use-case maps across risk, compliance, operations, and CX
- Reference architectures for credit risk AI and AML/KYC systems
- Model governance and documentation templates aligned to best practice
- Regulatory alignment checklists including EU AI Act high-risk requirements
- 12-month implementation roadmap with phase milestones
- 30-point AI governance readiness checklist

Contents

EXECUTIVE SUMMARY

Executive Summary

Financial services institutions operate at the intersection of regulation, risk, and customer trust. AI is already reshaping this landscape — from automating compliance and detecting fraud to accelerating credit decisions and personalising customer interactions. The question for 2026 is no longer whether to adopt AI, but how to do so in a governed, measurable, and regulatory-compliant way.

This blueprint translates the AI in Financial Services playbook into four practical tools: use-case maps, reference architectures, governance templates, and regulatory checklists. Use it alongside your institution's risk appetite, model risk management framework, and strategic plan.

$200–340B	83%	29%
Estimated annual value gen AI could add to global banking (McKinsey)	Reduction in AML regulatory penalties after AI-enhanced SAP rollouts	Improvement in insurance fraud detection with AI (Allianz Incognito)

Three Commitments for 2026

#	Commitment	What It Means in Practice
01	Use-Case Prioritisation	Identify 3–5 lighthouse AI initiatives with clear ROI and manageable regulatory exposure. Map each to value (P&L;, capital, cost, CX) and risk (model risk, conduct, regulatory).
02	Governed Architecture	Choose your AI architecture before engineering begins. Implement model registries, feature stores, and monitoring from day one. AI must interact with core systems via APIs, not direct database access.
03	Regulatory Alignment	Assess EU AI Act obligations for each use case. Maintain AI inventories with owners, risk ratings, and documentation. Build human-in-the-loop controls and explainability into all sensitive decisions.

How to use this blueprint: Read Parts 1–2 before your next architecture review. Use Part 3 templates as starting points for your model documentation. Complete the Part 4 checklists for each AI system under development. Revisit the roadmap quarterly.

PART 1 · BANKING & INSURANCE AI USE-CASE MAPS

1. Banking & Insurance AI Use-Case Maps

The tables below map AI use cases by domain and function. Maturity is rated: **Proven** (deployed at scale with published results), **Emerging** (early deployments, clear ROI evidence), or **Developing** (pilots underway, ROI not yet established at scale).

1.1 Risk & Compliance

Use Case	Domain	Maturity	Evidence & Impact
AML Transaction Monitoring	Operations	Proven	61% faster risk ID; 83% reduction in AML penalties. Real-time anomaly detection across millions of transactions.
KYC / Customer Due Diligence	Compliance	Proven	99.3% accuracy in identifying high-risk customers; 63% reduction in compliance officer workload.
Suspicious Activity Detection	Compliance	Proven	59% improvement in detection accuracy after AI-enhanced SAP rollouts.
Regulatory Reporting Automation	Compliance	Emerging	AI aggregates data, flags anomalies, and drafts regulatory reports. Reduces analyst time by 40–60%.
Credit Concentration Risk	Risk Management	Proven	41.7% improvement in risk measurement accuracy; 17.4% reduction in credit concentration exposure.
Fraud Detection (Card & Payments)	Operations	Proven	JPMorgan Chase: 20% reduction in validation rejections. Real-time pattern recognition beyond rule-based systems.
Model Risk Management	Risk Management	Emerging	AI-assisted model validation, back-testing, and performance monitoring. Reduces MRM team workload.
Conduct Risk Surveillance	Compliance	Developing	NLP on communications (email, chat, voice) to flag potential mis-selling and conduct breaches.

1.2 Credit & Customer Value

Use Case	Domain	Maturity	Evidence & Impact
Credit Early-Warning Systems	Credit Risk	Emerging	Summarises portfolio and borrower-level signals (behavioural, transactional, macro) to flag emerging risk earlier than traditional models.
Credit Memo Drafting	Credit Operations	Emerging	Gen AI turns structured and unstructured data (financials, news, management commentary) into first-draft memos for RMs.
Credit Decision Support	Credit Risk	Developing	Scenario analysis and rationale summaries alongside quantitative models. Adds interpretability without replacing models.

Use Case	Domain	Maturity	Evidence & Impact
Virtual Financial Assistants	Retail Banking CX	Proven	Answers queries on balances, fees, budgeting; provides personalised nudges. 75% of gen AI near-term banking value.
RM Copilots / Agent Assist	Corporate Banking	Emerging	Real-time answers, next-best actions, and document summarisation for contact centre and branch staff.
Personalised Cross-Sell	Retail Banking CX	Emerging	AI-driven cross-sell and upsell tailored to customer goals and risk appetite. Improves conversion and reduces churn.
Loan Origination Automation	Credit Operations	Proven	Document extraction, data validation, and decision-support automation. 40–60% faster origination cycle times.

PART 1 · BANKING & INSURANCE AI USE-CASE MAPS (CONTINUED)

1.3 Insurance: Underwriting, Claims & Fraud

Use Case	Domain	Maturity	Evidence & Impact
Submission Triage & Prioritisation	Underwriting	Proven	Automatically extracts key data, scores risk, and routes submissions. Reduces underwriter review time by 30–50%.
Misrepresentation & Ghost Broking Detection	Underwriting	Proven	Detects inconsistencies in applications before binding coverage. Shaves up to 5 points from combined ratios.
Coverage Analysis	Underwriting	Emerging	Reads policies, endorsements, and contracts to flag gaps or non-compliance with required coverage terms.
Automated FNOL Processing	Claims	Proven	Extracts data from forms, images, telematics, and third-party feeds. Reduces intake time by 60–80%.
Claims Triage & Routing	Claims	Proven	Risk-based prioritisation and assignment. 5–10× faster claim cycles; 20–50% reduction in resolution costs.
Damage Assessment (Computer Vision)	Claims	Emerging	AI assesses damage from photos and generates cost estimates. Reduces need for on-site assessments.
Claims Fraud Detection	Fraud	Proven	Cross-claim pattern analysis, provider/broker networks. 29% improvement in fraud detection (Allianz Incognito).
Customer Communication AI	Claims / CX	Emerging	AI assistants request documents, provide status updates, and explain decisions. Reduces inbound call volume.

1.4 Use-Case Prioritisation Matrix

Use the matrix below to prioritise your AI investments. Score each use case on value potential (1–5) and implementation feasibility (1–5). Lighthouse initiatives score ≥4 on both dimensions.

Use Case	Value (1–5)	Feasibility (1–5)	Recommendation
AML/KYC Automation	5	5	Start now — proven technology, strong vendor ecosystem, clear regulatory benefit.
Fraud Detection	5	4	High priority — strong ROI evidence, manageable model risk.
Credit Memo Drafting	4	5	Quick win — gen AI layer over existing data, low regulatory complexity.
Claims Triage (Insurance)	5	4	High priority — clear cost and cycle time ROI; proven at scale.
Early-Warning Systems	4	3	Medium priority — requires feature store and data infrastructure investment.
Virtual Financial Assistants	4	4	Accelerate — significant CX value, manageable risk with guardrails.
Conduct Risk Surveillance	3	2	Plan — valuable but high regulatory scrutiny and data complexity.

Use Case	Value (1–5)	Feasibility (1–5)	Recommendation
Coverage Analysis (Insurance)	3	3	Pilot — emerging technology, validate accuracy before scaling.

PART 2 · REFERENCE ARCHITECTURES FOR RISK & CREDIT AI

2. Reference Architectures for Risk & Credit AI

Effective FS AI architectures separate data, model, and application layers, with strong governance at each interface. The three architectures below cover the highest-priority domains: credit risk, AML/KYC compliance, and insurance claims. Each follows a layered pattern with model registries, feature stores, and API-based integration with core systems.

2.1 Credit Risk AI Architecture

Design principle: Combine traditional statistical models (PD, LGD, EAD) for regulatory capital with machine learning for early-warning and gen AI for summarisation and decision support. Each layer must be independently governed, documented, and validated.

Layer	Components	Governance Requirements
Data Layer	Core banking ledger, transaction history, bureau data, macroeconomic feeds, management accounts, news/NLP feeds	Data lineage tracking, PII controls, data quality SLAs, retention policies
Feature Store	Pre-computed borrower signals: behavioural scores, utilisation trends, covenant proximity, sector stress indicators	Feature versioning, documentation, drift monitoring, access controls
Model Layer (Statistical)	PD / LGD / EAD models (IRB or IFRS 9), scorecards, vintage analysis, stress testing models	SR 11-7 / SS1/23 validation, annual review, back-testing, MRM sign-off
Model Layer (ML)	Gradient boosting early-warning classifiers, anomaly detection for covenant breach prediction, SHAP explainability layer	Challenger/champion framework, bias assessment, explainability documentation
Gen AI Layer	LLM-based memo drafting, portfolio summary generation, borrower signal narration, RM copilot (meeting prep, proposals)	Prompt governance, output review workflow, hallucination controls, human sign-off
Application Layer	RM workbench, credit committee dashboard, early-warning alerts, customer-facing credit portal	Role-based access, audit trail, override logging, escalation workflow
Integration	Core banking via REST APIs, CRM write-back, data warehouse for reporting, regulatory reporting pipeline	API versioning, circuit breakers, no direct DB access

2.2 AML / KYC Compliance AI Architecture

Design principle: AI in AML/KYC must support — not replace — human judgement. Every AI-generated decision must be explainable, auditable, and subject to human review. The EU AI Act classifies AML/KYC as high-risk; full documentation and human oversight are mandatory.

Layer	Components	Governance Requirements
Data Ingestion	Transaction feeds, customer master data, sanctions/PEP lists, adverse media, corporate registry, correspondent bank data	Real-time and batch ingestion, data quality controls, source reconciliation
Entity Resolution	Graph database linking customers, accounts, beneficiaries, and counterparties. Network analysis for relationship mapping.	Deduplication rules, manual override workflow, data stewardship
Risk Scoring Engine	ML-based customer risk scoring (CDD/EDD), dynamic risk segmentation, periodic re-scoring triggers	Model documentation, validation, bias/fairness assessment, annual review
Transaction Monitoring	Real-time anomaly detection, typology-based rule engine, AI alert prioritisation and triage	Alert rationale logging, false positive tracking, tuning governance
SAR/STR Workflow	AI-drafted SAR narratives, investigator review queue, case management integration, regulatory submission pipeline	Human review mandatory before submission, four-eyes approval, audit trail
Regulatory Reporting	Automated report generation, threshold monitoring, cross-border reporting coordination	Reconciliation controls, legal review gates, version control
Observability	Model drift monitoring, alert volume dashboards, false positive/negative rates, SLA tracking	Daily monitoring, monthly model review, quarterly governance committee

PART 2 · REFERENCE ARCHITECTURES (CONTINUED)

2.3 Insurance Claims AI Architecture

Design principle: Claims AI should automate intake, triage, and routine decisions, while escalating complex, high-value, or disputed claims to experienced adjusters. Fraud detection must run across both policy and claims data for maximum efficacy.

Layer	Components	Governance Requirements
FNOL Intake	Multi-channel intake (web, app, phone, IoT/telematics), document and image extraction, OCR, structured data extraction	Data quality validation, completeness checks, PII handling, consent capture
Triage & Classification	ML-based claim severity scoring, coverage eligibility check, fraud risk flag, adjuster assignment	Model documentation, validation, bias testing (by claim type, geography)
Coverage Validation	Policy-to-claim mapping engine, endorsement and exclusion parsing, automated coverage determination	Audit trail for every determination, human override for edge cases
Damage Assessment	Computer vision for photo/video analysis, repair cost estimation, third-party valuation integration	Model accuracy benchmarking, human review for high-value claims
Fraud Detection	Cross-claim pattern analysis, provider/broker network analysis, anomaly scoring, link analysis across policy and claims	Investigation workflow, evidence packaging, litigation-ready audit trail
Decisioning & Payment	Automated settlement for fast-track claims, human review queue for complex/disputed claims, payment authorisation	Dual-control for settlements above threshold, customer notification log
Customer Communication	AI assistant for status updates, document requests, decision explanations; escalation to human agent	Communication logging, tone and accuracy review, complaint tracking

2.4 Architecture Selection Guide

Criterion	Single AI Model + Tools	AI Platform with Registries	Full Layered Architecture
Workflow complexity	Low (1–2 use cases)	Medium (3–6 use cases)	High (7+ use cases / multi-domain)
Regulatory exposure	Low	Medium	High (EU AI Act high-risk)
Implementation speed	Fastest (weeks)	Medium (months)	Slowest (quarters)

Criterion	Single AI Model + Tools	AI Platform with Registries	Full Layered Architecture
Observability	Basic logging	Model monitoring dashboard	Full governance layer
Best for	Pilots and PoCs	Department-level AI	Enterprise-wide AI programme
Minimum governance	Model card + owner	Registry + validation	Full MRM + AI Act compliance

PART 3 · MODEL GOVERNANCE & DOCUMENTATION TEMPLATES

3. Model Governance & Documentation Templates

The templates below provide starting-point structures for model documentation and risk assessment. Adapt them to your institution's model risk management policy, regulatory requirements, and internal governance framework. Every AI system that influences a material decision should have a completed model card.

3.1 Model Card Template

Field	Content Guidance	Why It Matters
Model Name	[Name] — [Version] — [Date]	Unique identifier and version tracking
Model Owner	[Name, Role, Business Unit]	Named individual accountable for the model
Model Purpose	One paragraph: what the model does, what decision it supports, and what it does not do.	Clear scope boundaries prevent scope creep
Training Data	Sources, date range, size, known limitations, PII handling	Required for EU AI Act high-risk documentation
Model Type	Statistical / ML / Gen AI / Ensemble	Determines applicable validation standards
Performance Metrics	Key metrics (AUC, F1, precision, recall, MSE as applicable) on train/test/OOT	Baseline for ongoing monitoring
Bias & Fairness	Protected characteristics tested, results, mitigation actions taken	Required for credit, pricing, claims decisions
Explainability	Explainability method (SHAP, LIME, etc.), example explanations, limitations	Required for EU AI Act high-risk; good practice for all
Human Oversight	Decision thresholds above which human review is mandatory; override process	Mandatory for high-risk AI under EU AI Act
Known Limitations	Conditions where model may underperform; out-of-scope scenarios	Supports appropriate use and escalation
Validation Status	Last validation date, validator, outcome, next scheduled review	Minimum annual review for material models
Risk Rating	Low / Medium / High / Critical — with rationale	Drives governance intensity and monitoring frequency
Change Log	Date, change made, approved by	Audit trail for model lifecycle management

PART 3 · MODEL GOVERNANCE & DOCUMENTATION TEMPLATES (CONTINUED)

3.2 Model Risk Assessment Framework

Rate each dimension 1 (low) to 5 (high). Sum the scores to determine the aggregate risk rating and required governance intensity. Any dimension scoring 5 triggers enhanced oversight regardless of aggregate.

Risk Dimension	Assessment Question	Scoring Guidance
Materiality	What is the financial, capital, or customer impact if the model fails or produces biased output?	1 = Minimal / 5 = Systemic
Regulatory Exposure	Is this a high-risk AI system under EU AI Act, EBA, PRA, or other FS-specific guidance?	1 = No regulatory overlay / 5 = High-risk AI Act
Data Quality	How well understood and controlled is the training and input data? Are there significant gaps or biases?	1 = Curated, audited / 5 = Unstructured, unvalidated
Model Complexity	Can the model's outputs be explained? Is it interpretable by domain experts without technical support?	1 = Simple scorecard / 5 = Large LLM, black box
Human Oversight	How much human review sits between model output and consequential action?	1 = Human decides / 5 = Fully automated, no review
Novelty	How established is the model type and vendor in this context? Is it a first-of-kind deployment?	1 = Established methodology / 5 = Novel or unproven
Operational Dependency	How many processes or decisions depend on this model running correctly?	1 = One team / 5 = Enterprise-critical dependency

Aggregate Score → Risk Rating

Aggregate Score	Risk Rating	Required Governance
7–14	Low	Standard model card; annual review; basic monitoring
15–21	Medium	Full model card; semi-annual validation; automated drift monitoring; MRM sign-off
22–28	High	Full model card + risk assessment; quarterly review; human-in-the-loop controls; senior sign-off
29–35	Critical	All of the above + independent model validation; Board/risk committee reporting; EU AI Act registration

3.3 Model Validation Checklist

Data Validation

- Training data sources documented and accessible

- Data lineage from source to model confirmed
- Training/test/OOT split documented and reproducible
- Known data quality issues logged and mitigated
- PII handling reviewed and approved by data privacy team

Model Performance

- Performancebenchmarked against predecessor or challenger model
- MetricsreportedonOOT and recent live data
- Sensitivityanalysiscompleted (key input ranges tested)
- Stresstestingcompleted for macroeconomic/tail scenarios
- Documentedthresholds for performance-triggered review

Bias & Fairness

- Protectedcharacteristics identified (age, gender, geography, etc.)
- Disparateimpact analysis completed on model outputs
- Mitigationactions documented where bias identified
- Post-deployment bias monitoring configured
- Legalandcompliance review completed

Explainability

- Explainabilitymethod selected and implemented (SHAP, LIME, or equivalent)
- Sampleexplanations reviewed for accuracy and clarity by domain expert
- Adverseactionreason codes generated and tested (for credit decisions)
- Modeldocumentation references explainability method and limitations
- Customer-facing explanation templates reviewed and approved

Governance

- Modelcardcompleted and approved by model owner
- Riskratingassigned using the Model Risk Assessment Framework
- Human-in-the-loop thresholds defined and tested
- Incidentresponse and rollback plan documented
- Modelversion control and change management process in place

PART 4 · REGULATORY ALIGNMENT CHECKLISTS

4. Regulatory Alignment Checklists

Financial services AI operates under a rapidly evolving regulatory landscape. The EU AI Act, effective from August 2024 with phased obligations through 2027, classifies many FS AI systems as high-risk. Non-compliance can result in fines up to €35 million or 7% of global annual turnover. Use the checklists below to assess and document your compliance posture.

4.1 EU AI Act — High-Risk FS Use Cases

The following FS AI use cases are likely to be classified as high-risk under the EU AI Act (Annex III). High-risk systems require conformity assessments, CE marking (where applicable), and registration in the EU database before deployment.

FS AI Use Case	Risk Category	Registration Required	Key Determination
Credit Scoring & Lending Decisions	High-Risk	Yes	Any AI that influences creditworthiness assessment or access to financial services
AML / KYC Customer Risk Scoring	High-Risk	Yes	AI systems used for anti-money laundering and know-your-customer risk classification
Insurance Underwriting Scoring	High-Risk	Assess	Risk classification of natural persons for health and life insurance underwriting
Claims Fraud Detection	High-Risk	Assess	Automated fraud decisions affecting claim outcomes for natural persons
Employment-Related HR AI	High-Risk	Yes	AI used in recruitment, performance management, or workforce decisions
Customer-Facing Virtual Assistants	Limited Risk	No	Chatbots interacting with customers — transparency obligations apply
Gen AI Tools (Copilots, Drafting)	Limited Risk	No	Must disclose AI-generated content; general-purpose AI model obligations apply
Internal Analytics / Dashboards	Minimal Risk	No	No specific EU AI Act obligations beyond general data protection compliance

4.2 EU AI Act Compliance Checklist — High-Risk Systems

For each AI system classified as high-risk, the following requirements must be met before deployment. Use this checklist during development and as part of your pre-deployment conformity assessment.

Risk Management System

- Documented risk management system covering the full AI lifecycle
- Risk identification and analysis for each high-risk use case
- Risk evaluation and mitigation measures documented

- Residual risk accepted and signed off by accountable owner
- Risk management system reviewed and updated at least annually

Data Governance

- Training,validation,and test data described and documented
- Dataqualitycriteriadefined and applied
- Examinationforbiases relevant to protected characteristics
- Datacollection,labelling, and preparation procedures documented
- Relevantdatagapsand shortcomings identified and addressed

Technical Documentation

- Purposeandintendeduseofthe AI system documented
- Architecture,components,and algorithms described
- Performancemetricsandbenchmarks reported
- Hardwareandsoftwarerequirements specified
- Changestothesystemduring lifecycle tracked and versioned

Transparency & Explainability

- Instructionsforuseprovidedtodeploying organisations
- Capabilities,limitations,andforeseeable misuse documented
- Levelofaccuracy,robustness,and cybersecurity specified
- Humanoversightmeasuresdescribed in instructions
- Explanationmechanismforindividual decisions implemented

Human Oversight

- Human-in-the-loopcontrols defined for consequential decisions
- Abilityforhumanoperators to understand system outputs
- Overridecapabilityimplemented and tested
- Monitoringofsystem operation during use configured
- Proceduresforsuspending or stopping the system documented

Accuracy, Robustness & Cybersecurity

- Accuracylevelsspecifiedandvalidatedforintended purpose
- Robustnesstoerrors,faults,orinconsistencies tested
- Resilience to adversarial inputs assessed
- Cybersecurity controls implemented and reviewed
- Performanceconsistencyacrossdiversepopulations tested

Registration & Conformity

- SystemregisteredinEUAIAct database (where required)
- Conformityassessmentcompleted or contracted
- Declaration of conformity prepared
- CEmarkingapplied(whereapplicable to product category)
- Post-market monitoring plan established

PART 4 · REGULATORY ALIGNMENT CHECKLISTS (CONTINUED)

4.3 FS AI Governance Checklist — Beyond EU AI Act

In addition to the EU AI Act, FS institutions must align with sector-specific guidance from the EBA (AI governance), PRA (model risk management), FCA (algorithmic accountability), and international standards (ISO 42001 AI Management System, NIST AI RMF).

AI Strategy & Governance

- Board-approved AI strategy linked to risk appetite and business priorities
- AI governance committee established with cross-functional membership
- AI governance integrated with existing model risk and operational risk frameworks
- Named AI ownership (strategy, architecture, compliance, operations)
- AI ethics principles documented and embedded in development processes

AI Inventory & Classification

- Complete inventory of AI and advanced analytics systems maintained
- Each system classified by risk level, purpose, and EU AI Act category
- System owners, users, and accountable executives identified
- Inventory reviewed and updated at least quarterly
- Regulatory reporting obligations mapped to inventory entries

Third-Party and Vendor AI

- Third-party AI systems included in AI inventory
- Contractual rights to audit, validate, and obtain documentation from vendors
- Vendor due diligence includes AI governance and EU AI Act compliance
- Model risk assessment completed for third-party models
- Exit strategy and data portability plan for each vendor dependency

Operational Controls

- Model monitoring dashboards operational for all material AI systems
- Automated alerts configured for performance degradation and drift
- Incident response and escalation procedures documented and tested
- Regular model validation cycle scheduled and resourced
- Business continuity plan covers AI system failures

Talent & Culture

- Model risk specialists in place or accessible (internal or advisory)
- AI governance leads identified across business units
- Risk and compliance staff trained on AI-specific risks
- Product owners for AI systems have domain and governance capability
- AI literacy programme in place for senior leadership and Board

PART 5 · 12-MONTH IMPLEMENTATION ROADMAP

5. 12-Month AI Implementation Roadmap

The roadmap below structures AI deployment across four phases, from baseline and prioritisation through to full industrialisation and regulatory alignment. Each phase has defined milestones, KPIs, and go/no-go criteria before proceeding to the next phase.

PHASE 1 · Months 0–3 — Baseline & Prioritisation

Milestone	Detail
Complete AI inventory	Document all current AI and advanced analytics use cases across risk, compliance, operations, and CX.
Value and risk mapping	Score each use case on value potential (P&L;, capital, cost, CX) and risk (model risk, conduct, regulatory). Use the prioritisation matrix in Part 1.
Select lighthouse initiatives	Identify 3–5 initiatives with ≥4/5 on both value and feasibility. Confirm with risk, compliance, legal, and IT.
Establish governance	Appoint AI governance lead. Establish AI governance committee. Align with MRM policy and model risk appetite.
Phase 1 KPIs	AI inventory complete; 3–5 initiatives selected; governance structure in place; regulatory assessment started.

PHASE 2 · Months 3–6 — Build Foundations & Pilot

Milestone	Detail
Strengthen model risk management	Ensure MRM function has AI-specific capability. Adopt model card template and risk assessment framework from Part 3.
Build core AI infrastructure	Implement or connect to model registry, feature store, and monitoring tooling. Ensure API-based integration pattern.
Launch pilots	Run one high-impact pilot in each of: risk/compliance (AML triage or KYC automation), credit (memo drafting or early-warning), insurance (claims triage).
EU AI Act pre-assessment	Complete use-case classification. Identify high-risk systems. Start documentation for conformity assessment.
Phase 2 KPIs	≥2 pilots live; model cards completed for pilot systems; monitoring dashboards operational; EU AI Act assessment in progress.

PHASE 3 · Months 6–9 **Scale & Integrate**

Milestone	Detail
Expand successful pilots	Scale pilots that met Phase 2 KPIs across products, geographies, and customer segments.
Integrate with operational systems	Connect AI outputs to core banking, policy admin, claims, and CRM systems via governed APIs.
Strengthen observability and controls	Deploy automated drift monitoring, bias/fairness checks, and performance dashboards for all live models.
Incident response readiness	Test rollback and suspension procedures for each AI system. Document escalation paths and accountabilities.
Phase 3 KPIs	≥4 use cases live at scale; all systems in model registry with validated model cards; zero unresolved compliance gaps.

PHASE 4 · Months 9–12 **Industrialise & Align with Regulation**

Milestone	Detail
Formalise AI inventories	Complete EU AI Act registration for high-risk systems. Finalise documentation and conformity assessments.
ISO 42001 / NIST AI RMF alignment	Conduct gap assessment against ISO 42001 and NIST AI RMF. Prioritise gaps. Begin certification process if applicable.
Embed AI metrics in business performance	Include AI KPIs (detection rates, cost savings, CX scores, model performance) in business scorecards and Board reporting.
AI talent and capability	Complete AI literacy programme for leadership. Confirm model risk and AI governance roles are resourced and permanent.
Phase 4 KPIs	Full EU AI Act compliance for high-risk systems; AI metrics in Board reporting; AI governance embedded in operational processes.

PART 6 · 30-POINT AI GOVERNANCE READINESS CHECKLIST

6. 30-Point AI Governance Readiness Checklist

Score your institution on each item: **Yes (2)** / **In Progress (1)** / **No (0)**. A score of ≥ 48 indicates strong governance readiness. Scores below 30 indicate significant gaps requiring immediate attention before scaling AI deployments.

Strategy & Governance (Items 1–6)	Yes	In Progress	No	Notes
1. AI strategy approved by Board or Executive Committee, linked to risk appetite				
2. AI governance committee established with cross-functional representation				
3. AI inventory maintained with owners, risk ratings, and documentation				
4. AI ethics principles documented and embedded in development processes				
5. Named AI accountability (strategy, architecture, compliance, operations)				
6. AI governance integrated with model risk and operational risk frameworks				

Risk & Compliance (Items 7–13)	Yes	In Progress	No	Notes
7. AI-enabled AML/KYC or transaction monitoring in place or actively piloted				
8. Documented improvements in detection accuracy and false-positive reduction				
9. Regulatory reporting and risk dashboards enhanced by AI summarisation				
10. EU AI Act or equivalent regulatory impact assessed for all AI systems				
11. High-risk AI systems identified and documentation programme underway				
12. Model risk management function has AI-specific capability and resources				
13. Third-party AI systems included in inventory and subject to vendor due diligence				

Credit & Customer Value (Items 14–18)	Yes	In Progress	No	Notes
14. Gen AI pilots in credit (early-warning, memo drafting, or customer engagement) underway				
15. Controls for explainability and bias for any AI-influenced credit decisions				
16. Virtual assistants or AI copilots supporting customer and RM interactions				
17. Adverse action reason codes generated and tested for AI-influenced credit decisions				
18. Customer consent and transparency obligations met for AI-driven personalisation				

Insurance Underwriting & Claims (Items 19–22)	Yes	In Progress	No	Notes
19. AI in underwriting submissions triage and risk assessment				
20. Claims AI for FNOL, triage, coverage validation, and decision support				
21. Fraud detection models deployed across both policy and claims				
22. Measured impact on combined ratio, loss ratio, and cycle times				

Architecture & Operations (Items 23–27)	Yes	In Progress	No	Notes
23. Layered AI architecture with model registry, feature stores, and monitoring				
24. APIs (not direct DB access) used for AI integration with core systems				
25. Logging, tracing, and incident response in place for all AI systems				
26. Regular model validation, back-testing, and performance reviews scheduled				
27. Business continuity and rollback plans cover AI system failures				

Talent & Culture (Items 28–30)	Yes	In Progress	No	Notes
28. Model risk specialists in place or accessible internally or via advisory				
29. Risk and compliance staff have completed AI-specific training				
30. AI literacy programme delivered or scheduled for senior leadership and Board				

Interpreting your score: 48–60: Strong governance readiness — focus on optimisation and regulatory alignment. 30–47: Moderate readiness — prioritise gaps in highest-risk areas. Below 30: Significant gaps — pause new AI deployments and address governance foundations first.

HOW FRACTO CAN HELP

How Fracto Can Help

Fracto works with banks, insurers, and financial services firms to turn AI strategy into governed, value-generating capability. We combine deep FS domain knowledge with AI architecture and governance expertise to help you move from isolated experiments to enterprise-scale deployment.

AI in FS Risk & Value Assessment

A structured engagement that identifies your highest-value AI opportunities, assesses risk and regulatory implications, and designs a 12-month roadmap for pilots, scaling, and controls. Delivered as a 2–3 week sprint with executive readout and actionable recommendations.

Architecture Design & Review

Independent review of your AI architecture against the reference patterns in this blueprint. Covers data layer, model stack, integration approach, observability, and governance controls. Delivered as a written assessment with prioritised recommendations.

Model Governance Programme

Design and implementation of your model governance framework: inventory, model cards, risk assessment, validation cycle, and monitoring. Aligned to EU AI Act, SR 11-7/SS1/23, and ISO 42001. Includes training for risk, compliance, and AI teams.

EU AI Act Compliance Readiness

Use-case classification, high-risk system identification, documentation programme, and conformity assessment support. Pragmatic, proportionate approach designed for FS institutions at any stage of AI maturity.

AI Pilot Acceleration

Hands-on support to design, build, and validate your first high-impact AI use case: AML triage, KYC automation, credit memo drafting, or claims triage. Includes use-case scoping, architecture, model governance, and stakeholder alignment.

Ongoing Advisory & Governance Support

Retained advisory covering AI governance committee support, regulatory monitoring, model review participation, and quarterly AI programme health checks. Flexible engagement model to complement your internal team.

Book an AI in FS Risk & Value Assessment fracto.ie · contact@fracto.ie