RJL SOLUTIONS

GOVERNANCE GUIDE / EVIDENCE AND OVERSIGHT

AI Governance Evidence

US bank and insurance expectations, the NIST AI RMF, and the UK and EU rules for international firms: what reviewers actually need to inspect.

SCOPE · CONTROLS · RELEASE · OPERATING EVIDENCE

Richard White
RJL Solutions · rjlsol.com

OCTOBER 2026

01 / THE PREMISE

Build the evidence around the deployed capability.

An enterprise should be able to show who approved an AI capability, what that approval covers and whether its controls work where people rely on them. Build that evidence into delivery and operation. A policy describes an expectation. A release record and an operating sample show whether anyone met it.

This guide is for CDOs, AI leaders, risk, compliance and model risk teams, CISOs and internal audit at US banks and insurers. A closing chapter covers firms that also operate in the UK or the European Union. The legal references identify the relevant questions. The implementation recommendations explain how to make the answers inspectable. Treat the questions as suggested preparation for a review; they aren't a universal auditor checklist.

BASIS	HOW TO USE IT
US supervisory and state rules	For banks: model risk, third-party and consumer-credit expectations. For insurers: state AI bulletins and regulations
NIST AI RMF	Organize risk management through a voluntary framework, and record the version used
UK and EU rules	For international firms: the EU AI Act, DORA and UK model risk principles, in sections 12 to 14
Enterprise control standard	Set accepted scope, tests, owners, release conditions and continuing service obligations

NIST describes AI RMF 1.0 as voluntary and currently under revision. This guide uses its published Govern, Map, Measure and Manage structure, together with the 2024 Generative AI Profile. A mapping to NIST does not establish legal compliance, or a NIST certification.[1, 2]

Make the unit of approval explicit

Identify the workflow and affected population, the model and configuration, the prompts, retrieval sources, semantic definitions, tools, permissions and user interface. Include the people and the downstream process that use the result. Model procurement evidence alone can't show this composition suits its intended purpose.

The recommended evidence package has a system record, an applicability assessment, a risk register, accepted contracts, evaluation results, a release decision and an operating history. Link each artifact to a stable system identifier and the version it describes. Record gaps and restrictions next to the favorable results, so the approved scope stays visible.

WHAT THIS GUIDE IS NOT	A delivery and evidence guide, not a legal opinion or a conformity assessment. Legal and compliance owners validate applicability and duties against current law. References reviewed 10 October 2026.

02 / US BANKS

US banks: model risk, third parties and credit decisions.

For US banks, the expectations that govern an AI capability come from supervisory guidance and consumer law. No single AI statute covers them. Record which apply to each capability, and who in the second line reviews it.

EXPECTATION	WHAT IT ASKS	EVIDENCE FROM THIS GUIDE
SR 26-2, Revised Guidance on Model Risk Management (Federal Reserve, OCC and FDIC, 17 April 2026)[21]	Supersedes SR 11-7 and SR 21-8. Nonbinding; most relevant above $30 billion in assets, with smaller-bank exceptions for significant model risk. Footnote 3 excludes generative and agentic AI	Record the internal policy that governs each generative or agentic capability; the system record, evaluation and release decision are its evidence
Interagency Guidance on Third-Party Relationships (June 2023)[26]	Risk-based planning, due diligence, contract terms, ongoing monitoring and termination, including for model and AI providers	Supplier records, tested integration and an exit plan
ECOA and Regulation B, adverse action notices (12 CFR 1002.9)[27]	Specific principal reasons for a credit denial or other adverse action, however complex the model	Reasons traceable to the request trace, with tested explanations
NYDFS Part 500, with the October 2024 industry letter on AI[29]	AI risks addressed within existing risk assessment, third-party, access, training, monitoring and data controls; the letter adds no new requirements	Threat model, access and revocation tests, and monitoring of unusual queries

The CFPB's 2022 circular on adverse action notices for credit decisions based on complex algorithms was withdrawn in May 2025, with dozens of other guidance documents. The Regulation B duty to give specific reasons was not.[28]

Because SR 26-2 leaves generative and agentic AI out of scope, a bank can run a mature model risk program and still have no governing route for an AI assistant. Close that gap explicitly. Name the policy, the reviewer and the evidence standard for each capability.

ASK EARLY Which framework governs this capability, and who in the second line will review it? The answer usually sets the calendar.

03 / US INSURERS

US insurers: governance programs and unfair discrimination testing.

States set insurance regulation of AI one by one. The common thread is a written governance program, oversight of vendors and external data, and testing for unfair discrimination in consumer outcomes.

EXPECTATION	WHAT IT ASKS	EVIDENCE FROM THIS GUIDE
NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers (December 2023)[22]	A written AI systems program covering governance, risk management, internal controls, third-party oversight and testing for adverse consumer outcomes. Adopted by many US insurance jurisdictions; check the NAIC's adoption map	Applicability record, supplier records, evaluation results and release decisions as program evidence
New York: NYDFS Insurance Circular Letter No. 7 (July 2024)[23]	Governance, risk management and fairness testing for AI systems and external consumer data used in underwriting and pricing	Data source approvals, and evaluation sets and release criteria that include fairness tests
Colorado: Regulation 10-1-1, for life insurers and, since 15 October 2025, private auto and health benefit plan insurers[30]	A risk-based governance framework, an inventory of external data, algorithms and models, board oversight, vendor oversight and annual compliance reports	System record and inventory, owner and board reporting, and supplier records

Several large states, including New York and Colorado, have their own rules instead of the bulletin. Check each state where the insurer writes business, and record the result in the applicability record.

Claims, underwriting and pricing workflows carry different exposure. Record which consumer decisions the capability influences, and test for unfair discrimination wherever it can affect an outcome. An internal claims assistant that only assembles case context is a lower-exposure place to start than a pricing model.

PROGRAM, NOT PROJECT Regulators expect an insurer-wide AI program. Each capability's record should show how it fits that program.

04 / RISK WORK

Organize risk work with the NIST AI RMF.

The AI RMF provides a common language for risk work across the lifecycle. Its four functions are not a one-time sequence, and Govern runs through the other three. Use the framework to connect business context, evaluation, decisions and continuing response. The mapping below is a recommended artifact pattern, rather than a NIST-prescribed document set.[1]

FUNCTION	PRACTICAL RECORD, AND THE ACCOUNTABLE DECISION
Govern	Policy, inventory, roles, risk authority, supplier rules and review cadence. Leadership establishes who may accept which risk.
Map	Intended use, affected groups, dependencies, foreseeable misuse and alternatives. Business and risk owners agree on the context.
Measure	Evaluation design, results, uncertainty and limitations. Qualified reviewers assess the relevant risks and performance.
Manage	Mitigation, release restrictions, monitoring, incident response and retirement. Authorized owners decide and track action.

Keep the adopted framework version in the enterprise policy. NIST notes that AI RMF 1.0 is being revised. A new publication should trigger an assessment of the mapping, rather than silently changing what a previous release was evaluated against.[2]

Adapt the profile to the actual risks

The Generative AI Profile identifies concerns such as confabulation, data privacy, information security, human-AI configuration and value-chain integration. Select the risks relevant to the workflow and connect them to evaluated controls. Its suggested actions include checking sources and citations, documenting limits of generalization, and using empirical evaluations rather than anecdotal demonstrations.[13]

For a claims context assistant, check three things. Is each factual statement supported by the permitted source? Is important contradictory information left out? Can a user verify the source? Assess user overreliance as well as answer quality. Explain how the interface tells a retrieved fact from a generated summary and from missing information.

Define risk acceptance as an operating decision

A residual-risk entry should describe the failure mode, the people affected, the remaining exposure, the controls and the uncertainty. Name the person with authority to accept it, the approved scope, the review date and the conditions that cancel the acceptance.

ON INDEPENDENCE — A delivery lead’s acknowledgment can’t stand in for a risk decision outside their authority. Match independence to the consequence of the use. The team that builds a capability shouldn’t be the only voice judging whether people can rely on it.

05 / THE CROSSWALK

Connect requirements to controls, and controls to evidence.

Give each requirement an identifier. Record its source, applicability, control owner, implementation, test method, evidence location and current status. Treat this register as a route into existing engineering and risk records. Copying the same narrative into several spreadsheets makes reconciliation harder and the evidence no better.

The examples below connect selected requirements across US and international regimes to recommended delivery artifacts. They're a partial working crosswalk. They don't state every obligation. A legal owner has to confirm which entries apply, and add the requirements left out for the system.

REQUIREMENT AREA AND SOURCES	EVIDENCE TO PREPARE FOR REVIEW
Risk management and validation: SR 26-2, NAIC bulletin, AI Act Art. 9[21, 22, 6]	Hazard and misuse register, independent review, mitigation tests, residual risks and an approved acceptance rationale
Data governance and fairness: NYDFS Circular Letter 7, Colorado 10-1-1, AI Act Art. 10[23, 30, 7]	Data source approvals, provenance, quality and unfair-discrimination testing, and limitations
Documentation: SR 26-2, AI Act Art. 11[21, 8]	A versioned system record and technical documentation
Record keeping: NYDFS Part 500 audit trail, AI Act Art. 12[29, 9]	Logging design and operating samples that demonstrate event traceability
Explanation of credit decisions: Regulation B[27]	Specific reasons tied to the request trace, and tested explanation quality
Human oversight: NAIC bulletin, AI Act Art. 14[22, 10]	Interface and workflow tests showing that competent people can interpret and intervene
Robustness and security: NYDFS Part 500, AI Act Art. 15[29, 11]	Task metrics, failure and attack tests, limitations, and operating observations

Show design and operating effectiveness separately

For an entitlement control, the design evidence might be the reviewed policy, the delegated identity contract and the enforcement architecture. Operating evidence is a sample of real allowed and denied requests, backed by negative tests of the deployed route. A diagram alone doesn't show that a cache, an index or a tool applies the same restrictions.

For each test, identify the approved criterion, environment, implementation version, input population, observed outcome, evaluator and date. An evidence link has to resolve for an authorized reviewer and identify the artifact it points to. Link failures and exceptions to the decisions that corrected or restricted them.

06 / RECONSTRUCTION

Reconstruct an authorized request.

An authorized reviewer should be able to follow a request from identity to result and use. Carry a correlation identifier through the application, agent, retrieval, data product and tools. Keep explanatory evidence without building an uncontrolled store of sensitive information.

Article 12 requires appropriate automatic logging capabilities for high-risk systems. The pattern below is an enterprise trace pattern. It is broader than the article's general wording, and must be tailored to purpose and applicable duties.[9]

TRACE ELEMENT	THE QUESTION THE EVIDENCE SHOULD ANSWER
Request context	Which human, application or delegated agent requested what operation, and for what purpose
Authorization	Which policy and identity decision allowed, masked, restricted or denied the request
Configuration	Which model, prompt, route, product, semantic, tool and policy versions were active
Information used	Which source records or document versions supported the answer, and what freshness limits applied
Outcome and intervention	What was returned, refused, changed, escalated or overridden, and by whom
Service observation	What errors, retries, timing and cost occurred, and which dependencies were involved

Protect the evidence itself

Set access, integrity, retention, deletion and export rules by evidence class. Use controlled references and redaction where they keep the explanation intact. A hash supports integrity checking. It can't reconstruct the source or prove lawful access. Source retention and deletion affect what you can reconstruct.

For applicable high-risk deployer logs under its control, Article 26 states a minimum retention period of six months unless applicable Union or national law provides otherwise. Do not apply that number indiscriminately to all prompts, documents, system records or personal data. Financial institutions must keep them as part of their financial-services documentation (Article 26(6)). Assign a legally approved retention schedule by record type.[14]

TEST RECONSTRUCTION FROM PRODUCTION	Sample a permitted request, a denial, an ambiguous answer and a human intervention. Recover the configuration, source context and authorization decision. Exact model replay may differ even with fixed settings, so keep the historical record of execution and the approved system state. A replay can supplement that record. It can't replace what happened.

07 / EVALUATION

Evaluate the workflow before release.

Evaluate the full capability in the conditions under which people will use it. Include representative work, foreseeable misuse, meaningful exceptions and critical control failures. Set the acceptance criteria before reviewing the result. Article 9 links high-risk testing to intended purpose and prior defined metrics and thresholds; Article 15 addresses accuracy, robustness and cybersecurity.[6, 11]

Use a test population with a stated boundary

Describe how cases were selected, the observation period, the affected population, labels and adjudication, and known coverage gaps. Separate development cases from a suitably protected evaluation set. Include difficult and uncommon cases where their consequences matter. Record uncertainty and sample limits. An average presented as a guarantee misleads.

TEST AREA	CONCRETE EVIDENCE FOR THE EXAMPLE ASSISTANT
Useful completion	The examiner receives the required context, with supported statements and acceptable verification effort
Source and meaning	Missing, stale, contradictory and semantically incompatible records produce the agreed handling
Access and isolation	Unauthorized claim, revoked user, cross-tenant request and restricted cached content are blocked
Adversarial behavior	Instructions in retrieved documents cannot authorize a tool or disclose restricted information
Failure and recovery	Timeouts, index lag, model changes and policy failures cause the accepted fallback or restriction
Human interaction	Users identify limitations, reject unsuitable output, and escalate without losing the evidence trail

Record results so another reviewer can assess them

Keep the test pack version, system version, environment, case count, metric definitions, denominators, raw outcomes, adjudication record, evaluator and execution time. Distinguish offline evaluation from observed production behavior. A test run against an older prompt or policy needs an explicit argument for why it still applies to the proposed release.

Report material errors separately by consequence and affected group, and include the effort needed to verify or correct outputs. Where lawful and appropriate, use relevant subgroup analysis to look for uneven performance, with privacy and legal owners governing the data it uses.

08 / OVERSIGHT

Give people the authority and the capacity to intervene.

For high-risk systems, Article 14 addresses effective oversight, and Article 26 requires deployers to assign it to people with competence, training, authority and support. Demonstrate that they can recognize a problem and take the intended action within the workflow.[10, 14]

ROLE	DECISION AND EVIDENCE RESPONSIBILITY
Business process owner	Approves intended use, workflow consequence, useful outcome, and human responsibility
Data and semantic owners	Approve source authority, meaning, permitted joins, quality limits, and source corrections
Legal, privacy and compliance	Decide applicability, required assessments, disclosure and retention duties, and legal escalation
Security and risk owners	Challenge controls, and authorize risk acceptance within their delegated authority
Service and product owners	Maintain tests, versions, monitoring, support, change control, and containment
Internal audit	Independently assesses relevant governance and control evidence within its mandate

Assign these responsibilities with funded capacity for review and exceptions. One person may carry several roles. Business approval, specialist challenge and independent assurance should still stay visible.

Demonstrate oversight in a realistic exercise

Test a misleading answer, a stale source and an unauthorized request in the actual interface, under representative workload. Watch verification, correction, rejection and escalation. Record gaps in instructions, permissions or available time. Provide a way to pause or restrict the capability and handle the underlying work safely. Train for scope, uncertainty, verification, automation bias and escalation, and assess competence, not just attendance.

Check impact and transparency duties

Article 27 requires a fundamental rights impact assessment for specified deployers and uses, including relevant creditworthiness and life or health insurance uses. Assess applicability rather than requiring it for every AI project. It can complement a data protection impact assessment, and the records should preserve the distinct questions and required actions.[15] Article 50 sets transparency duties for specified systems and content; legal owners determine the required notice or disclosure. Test its placement, and retain the approved wording and interface evidence.[16]

09 / SUPPLIERS

Keep supplier evidence tied to your system.

Ask suppliers for evidence that answers the questions your deployment creates. A model card, a security report or a statement of compliance can inform the assessment, but its scope may exclude your prompts, data, tools, delegated access and business process. Document which risks the supplier covers, and which remain with the enterprise. US banks should apply the 2023 interagency guidance on third-party relationships, and EU financial entities must also meet DORA's contract and register requirements.[26, 25]

SUPPLIER QUESTION	RECORD TO REQUEST OR ESTABLISH
Approved use and limits	Intended use, known limitations, relevant evaluation conditions, and prohibited uses
Data and processing	Processing locations, subprocessors, retention, training use, deletion, and contractual restrictions
Version and change	Model or service version, update notice, compatibility evidence, and the ability to test before reliance
Control and incident	Security responsibilities, incident notification route, evidence access, and cooperation obligations
Continuity and exit	Service dependency, recovery behavior, export or migration plan, and practical replacement constraints

Use contractual commitments and technical verification together. A supplier may not provide every artifact you ask for. Record the resulting uncertainty, and whether it's acceptable for the approved scope. Procurement and risk owners should resolve missing evidence before a dependency becomes essential to a consequential workflow.

Verify the dependency in your environment

Test the integration. Don't assume the supplier's evaluation describes your service. Check identity propagation, processing boundaries, error handling, update behavior and the ability to keep evidence. Exercise the agreed response to an outage, and confirm the business process can carry on safely, or stop within its accepted conditions.

Keep the supplier's evidence version, assessment date, reviewer, open questions and accepted restrictions with the system record. Give the product and procurement owners a review trigger for contract renewal, service changes and emerging incidents.

WHERE INSPECTION IS PREVENTED	If a supplier prevents meaningful inspection, the approval decision should state the limitation and its consequence.

Approve the release, and its continuing conditions.

The release decision should identify the system and environment, intended use, configuration and dependency versions, applicable controls, evaluation results, open issues, approval holders and operating conditions. Include the authorized user population, allowed actions, data boundary, support route and next review trigger.

Keep three decisions separate: funding a prototype, allowing production reliance and accepting operational transfer. Governance review must never widen authority by implication. If the approved scope is read-only assistance, write access or an automated business decision needs a new review of authority and consequence.

Make gaps and exceptions actionable

State whether each control is tested, has unresolved failures, is planned, is blocked or doesn't apply, with a rationale. For every material gap, show the consequence, owner and closure condition.

> **A favorable average score cannot compensate for unauthorized disclosure.**

Block release when access, isolation, action authority, integrity or required evidence can't be enforced. Fix the cause and repeat the relevant tests. For an exception policy permits, identify the accepting authority, exposure, compensating control, expiry and enforced restriction. No universal accuracy percentage proves suitability across workflows.

Reassess the complete composition after change

Assess model updates, prompts, retrieval corpora, indexes, semantic mappings, tools, permissions, source contracts and workflow changes for their effect on accepted behavior. Require tests in proportion to the change and the risk. When you test more narrowly, record why that's enough. Emergency changes need traceable approval and a follow-up review.

For US banks, SR 26-2's independent validation expectations apply to the models within its scope; for a generative or agentic capability outside it, the release record should name the reviewer who plays that role.[21] Article 25 makes value-chain role changes important, and Article 43 addresses conformity assessment for high-risk systems, including relevant modifications. Obtain legal review where a change could alter those duties. An internal version increment alone does not resolve a substantial-modification question.[5, 12] Legal conformity work may require additional documentation, quality management, registration, declarations, or a specific assessment procedure; an internal release review does not replace the applicable conformity route.[12]

Operate the evidence, and prepare for review.

Assign service owners to watch actual use and control behavior after release. Link incidents, complaints, overrides, source changes and control exceptions to the same system identifier. Reviews should show what changed since approval, what the observations imply and who authorized the response.

For applicable high-risk systems, Article 72 addresses provider post-market monitoring, Article 26 deployer monitoring, and Article 73 serious incident reporting. Map the relevant actor, reporting channel and deadline to a runbook. Clocks differ by regime: AI Act Article 73 serious incident reports, GDPR Article 33 breach notification within 72 hours, DORA major ICT incident reporting, and NYDFS Part 500.17 notice within 72 hours. Use the applicable requirement rather than one generic target.[14, 17, 18]

REVIEW QUESTION	EVIDENCE TO OPEN DURING THE REVIEW
What exactly was approved	Intended-use record, applicability decision, release scope, and configuration
Who could change or stop it	Delegated authority, access records, change approvals, and an intervention exercise
Did the controls work	Test failures as well as passes, and samples of operating allowed and denied requests
What harm or failure occurred	Incident and complaint records, impact assessment, containment, and correction
Why does continued use remain acceptable	Current observations, residual-risk review, open exceptions, and the owner decision

The reviewer decides the actual scope and criteria. Give them a population of records to sample from. A handful of favorable examples won't do.

Start with one inspectable capability

Choose a consequential repeated workflow, name its owner and find its current release decision. Trace one request, one denial, one failure and one human intervention. List the evidence gaps and what they mean. Set closure owners, and restrict the scope wherever missing evidence prevents justified reliance.

Frame, Prove, Build and Run & Transfer can make these records part of delivery: establish scope and applicability in Frame, test a complete slice in Prove, implement accepted contracts and controls in Build, and demonstrate continuing ownership through Run & Transfer. Define the engagement scope, specialist responsibilities, acceptance criteria and transfer conditions before work begins. Request a free sixty-minute working session at **rjlsol.com/contact** or write to **rwhite@rjlsol.com**.

Classify the intended use, and the legal role.

Start with the actual decision or assistance the system provides. Document where it operates, who it affects and how its output shapes the process. Have legal and compliance owners approve the assessment and its rationale. Reassess when the workflow, audience, authority or purpose changes.

Under Article 6 and Annex III, high-risk classification depends on specified uses and conditions. Annex III includes creditworthiness assessment of natural persons, with an exception for financial fraud detection, and risk assessment and pricing for natural persons in life and health insurance. This does not classify every banking or insurance assistant as high-risk, though an Annex III system that profiles natural persons always is.[3]

Article 6(3) lets a listed system fall outside high-risk when it performs a narrow procedural task, improves the result of a completed human activity, detects decision patterns without replacing human assessment, or performs a preparatory task. A provider relying on that exemption must document the assessment before placing the system on the market, and register it (Article 6(4)). Under Article 111(2), high-risk systems already on the market before the high-risk rules apply are covered only if their design changes significantly afterward, so treat significant changes to existing credit or pricing systems as a classification trigger.[3, 20]

> **Calling it read-only, or keeping a person in the workflow, does not settle classification.**

Annex III doesn't list claims handling itself. This guide uses an internal claims context assistant as its example and applies the high-risk requirements as a demanding reference. Reassess if its outputs feed underwriting, pricing or the profiling of natural persons.

Establish who carries which duty

Record whether the enterprise is a provider, deployer or other actor for each system. Article 25 identifies circumstances in which rebranding, substantial modification or changing intended purpose can move a party into provider responsibilities. Buying a model and deploying a system are distinct positions.[5]

THE APPLICABILITY RECORD	Names its approver, assessment date, source version, next review trigger and open questions.

Record the current application timetable.

The Digital Omnibus on AI, Regulation (EU) 2026/1744, entered into force on 27 July 2026 and moved the high-risk dates. Record the dates that apply to each system, and the duties that already apply, in the applicability record.[4, 19]

PROVISION GROUP	APPLICATION DATE
AI literacy (Article 4) and the original prohibited practices	2 February 2025
Governance and GPAI model obligations	2 August 2025
General application and transparency rules	2 August 2026, with specified exceptions
Article 50(2) marking, for generative systems placed on the market before 2 August 2026	2 December 2026
New Article 5 prohibitions on non-consensual intimate imagery and child sexual abuse material	2 December 2026
Annex III high-risk rules	2 December 2027
High-risk systems in regulated products	2 August 2028

Article 111(4), added by the Omnibus, gives providers of generative systems placed on the market before 2 August 2026 until 2 December 2026 to comply with Article 50(2).[20] The Omnibus also rewrote Article 4: providers and deployers must take measures to support the AI literacy of their staff, rather than ensure a set level for each person.

Provisions written for financial institutions

Several duties are met through existing financial-services rules. A deployer that is a financial institution satisfies its monitoring duty by complying with its internal governance requirements (Article 26(5)), and must keep the system logs as part of the documentation those rules require (Article 26(6)).[14] For the Annex III areas that include credit and insurance, conformity assessment follows internal control, without a notified body (Article 43(2)).[12]

EARLIER DATES STILL APPLY — A later high-risk date does not excuse duties that already apply, such as AI literacy, the prohibitions and the transparency rules.

Add the UK and EU financial-sector rules.

Firms with UK or EU operations add these to the US evidence. The same system record, supplier records, release decisions and incident runbook can serve them. Legal and compliance owners decide which apply.

REGIME	WHAT IT EXPECTS	HOW THIS EVIDENCE SERVES IT
UK banks: PRA SS1/23, Model Risk Management Principles (effective 17 May 2024)[24]	Model inventory, tiering, independent validation and governance, including AI-based models	System record, version evidence and independent review
EU financial entities: DORA, Regulation (EU) 2022/2554 (applies from 17 January 2025)[25]	ICT third-party risk, a register of information, contract terms and major ICT incident reporting	Supplier records, exit plans and the incident runbook
Personal data: GDPR Articles 22 and 35	Limits on solely automated decisions, and data protection impact assessments	Oversight design and the impact assessment record

For UK banks, the PRA's principles bring AI-based models into model risk management, unlike the US guidance. Keep one system record that answers both, and note where the scope differs.

ONE RECORD, SEVERAL REGIMES	Build the evidence once, around the deployed capability, and map it to each regime that applies. Don't run a separate program for every regulator.

ARCHITECTURE UPDATE / WORKED EVIDENCE CHAIN

Reconstruct the decision and the action.

Illustrative invoice request: explain a hold, then propose a permitted correction.

EVIDENCE LINK	MINIMUM RECORD FOR AN AUTHORIZED REVIEWER
Identity and context	Request ID, verified subject/workload references, tenant, task, context version and relevant freshness checks. Avoid copying unnecessary employee or customer details.
Read authorization	Policy version and decision reference, evaluated resource scope and permitted source versions. Preserve a denial record without leaking the denied content.
Generation	Model, prompt and tool versions; protected references to retrieved evidence; delivered response and limitations. Do not require hidden model reasoning.
Approval	Exact proposed payload or protected digest plus retained payload reference, authorized approver, timestamp, scope and expiry. A changed payload invalidates approval.
Execution	Fresh authorization and state checks, idempotency key, command result and external transaction reference. An uncertain timeout enters reconciliation before retry.
Operation	Exceptions, reviewer overrides, incident links and recovery decision. Evidence access, integrity and retention have named owners.

Test an allowed read, a denied cross-tenant read, revoked access between approval and execution, a stale invoice version and a timeout after the target accepted the command. The expected outcome is a specific response or escalation, with reconstructable evidence.

For a claims recommendation, add the relevant decision criteria, evaluation evidence and authorized reviewer disposition. This illustrative technical record supplements the applicable governance process; it does not settle legal classification or demonstrate compliance by itself.

Design references checked 29 October 2026: Data mesh principles · Document-level access control · RAG security. These additions are RJL implementation recommendations, not a vendor certification or a compliance determination.

Official sources and review basis.

Prepared by Richard White, 6 October 2026; not legal advice. Revised 10 October 2026. Next review by 2 December 2026, or sooner if the timetable changes. Service Desk pages are as published after the Omnibus; the Official Journal text governs. This edition summarizes the changes most relevant to banking and insurance evidence, not all of them; take legal advice.

US BANKING AND INSURANCE

[21] **SR 26-2.** Revised Guidance on Model Risk Management, Federal Reserve, OCC and FDIC, 17 April 2026; issued by the OCC as Bulletin 2026-13.
federalreserve.gov/supervisionreg/srletters/SR2602.htm

[26] **Interagency Guidance on Third-Party Relationships: Risk Management.** Federal Reserve, FDIC and OCC, June 2023.
federalregister.gov/d/2023-12340

[27] **Regulation B, 12 CFR 1002.9.** Notifications, including the statement of specific reasons for adverse action.
ecfr.gov/current/title-12/part-1002

[28] **CFPB withdrawal of guidance documents.** Includes Circulars 2022-03 and 2023-03, 12 May 2025.
consumerfinance.gov/compliance/guidance/withdrawn-guidance/

[29] **NYDFS industry letter.** Cybersecurity risks arising from artificial intelligence, 16 October 2024; 23 NYCRR Part 500.
dfs.ny.gov/…

[22] **NAIC Model Bulletin.** Use of Artificial Intelligence Systems by Insurers, December 2023.
content.naic.org/…

[23] **NYDFS Insurance Circular Letter No. 7.** Use of AI systems and external consumer data, July 2024.
dfs.ny.gov/industry-guidance/circular-letters/cl2024-07

[30] **Colorado Regulation 10-1-1.** Governance and risk management framework requirements for life insurers, effective 14 November 2023; amended to cover private passenger auto and health benefit plan insurers, effective 15 October 2025.
doi.colorado.gov/…

FRAMEWORK AND CURRENT TIMETABLE

[1] **NIST AI RMF Core.** Govern, Map, Measure, Manage, and the published categories.
airc.nist.gov/airmf-resources/airmf/5-sec-core/

[2] **NIST AI RMF overview.** Voluntary status, the 1.0 revision, and the companion profile.
nist.gov/itl/ai-risk-management-framework

[4] **European Commission AI Act overview.** Application dates and Omnibus changes.
digital-strategy.ec.europa.eu/…

[13] **NIST AI 600-1.** Generative Artificial Intelligence Profile, July 2024.
nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

[19] **Regulation (EU) 2026/1744.** The Digital Omnibus on AI.
eur-lex.europa.eu/eli/reg/2026/1744/oj

Requirements, operation and other regimes.

CLASSIFICATION AND ACTORS

[3] **AI Act Article 6 and Annex III.** Classification conditions, and the specified creditworthiness and insurance uses.
ai-act-service-desk.ec.europa.eu/en/ai-act/article-6 · ai-act-service-desk.ec.europa.eu/en/ai-act/annex-3

[5] **AI Act Article 25.** Responsibilities along the AI value chain.
ai-act-service-desk.ec.europa.eu/en/ai-act/article-25

SELECTED SYSTEM REQUIREMENTS

[6–8] **Articles 9, 10 and 11.** Risk management; data and data governance; technical documentation and Annex IV.
ai-act-service-desk.ec.europa.eu/en/ai-act/article-9 · ai-act-service-desk.ec.europa.eu/en/ai-act/article-10 · ai-act-service-desk.ec.europa.eu/en/ai-act/article-11 · ai-act-service-desk.ec.europa.eu/en/ai-act/annex-4

[9–11] **Articles 12, 14 and 15.** Record keeping; human oversight; accuracy, robustness and cybersecurity.
ai-act-service-desk.ec.europa.eu/en/ai-act/article-12 · ai-act-service-desk.ec.europa.eu/en/ai-act/article-14 · ai-act-service-desk.ec.europa.eu/en/ai-act/article-15

[12] **AI Act Article 43.** Conformity assessment.
ai-act-service-desk.ec.europa.eu/en/ai-act/article-43

DEPLOYMENT AND CONTINUING OPERATION

[14–16] **Articles 26, 27 and 50.** Deployer obligations; fundamental rights impact assessment; transparency obligations.
ai-act-service-desk.ec.europa.eu/en/ai-act/article-26 · ai-act-service-desk.ec.europa.eu/en/ai-act/article-27 · ai-act-service-desk.ec.europa.eu/en/ai-act/article-50

[17, 18] **Articles 72 and 73.** Post-market monitoring; serious incident reporting.
ai-act-service-desk.ec.europa.eu/en/ai-act/article-72 · ai-act-service-desk.ec.europa.eu/en/ai-act/article-73

[20] **AI Act Article 111.** Systems already placed on the market, including paragraph 4 added by the Omnibus.
ai-act-service-desk.ec.europa.eu/en/ai-act/article-111

UK AND EU FINANCIAL RULES

[24] **PRA SS1/23.** Model risk management principles for banks, May 2023, effective 17 May 2024.
bankofengland.co.uk/…

[25] **DORA.** Regulation (EU) 2022/2554 on digital operational resilience, applying from 17 January 2025.
eur-lex.europa.eu/eli/reg/2022/2554/oj

Every source is linked. Links updated 15 October 2026.

ABOUT

Richard White

I was global firm chief data officer at McKinsey & Company, owning data, AI and platforms end to end, and before that chief data officer at New York Life, where I built the chief data office and reported data risk to risk committees at multiple levels. Earlier I led data at the Federal Reserve Bank of New York and enterprise data governance at Citigroup. I founded RJL Solutions in 2024. My book, *The Day After the Demo*, is free at rjlsol.com/book.

RJL Solutions connects strategy, architecture, engineering, governance and operation through Frame, Prove, Build, and Run & Transfer. Each phase is separately scoped and priced in a written proposal before you commit. The first sixty-minute working session is free.

STANDARD **Data Product Certification Standard**
The input contracts this evidence depends on.

TECHNICAL **Agent Onboarding: Treating an Agent Like a Data Product**
Tool and agent composition, in 13 pages.

DELIVERY **The 90-Day First Governed Capability**
The delivery gates that produce these records, in 15 pages.

RJL SOLUTIONS

WEB rjlsol.com

EMAIL rwhite@rjlsol.com

This is a delivery and evidence guide. It isn't a legal opinion, a conformity assessment or a statement of every applicable duty. Legal and compliance owners have to validate applicability, transitions, sector duties and the required conformity route against current law. I reviewed the references on 10 October 2026; recheck them before use.