RJL SOLUTIONS

FLAGSHIP PAPER / RISK, LEGAL AND CYBER

Governing AI from Registry to Runtime

An integrated risk, legal and cyber control framework for banks and insurers: one AI product registry, one risk assessment that gates production, guardrails that enforce what was approved, and evidence from every decision.

REGISTRY · GATE · RUNTIME · EVIDENCE · THREE LINES

Richard White
RJL Solutions · rjlsol.com

VERSION 1.1
20 NOVEMBER 2026

IN BRIEF

Governing AI from Registry to Runtime

AI governance fails when risk, legal and security run separate processes: three reviews of three descriptions of the same system, on three calendars, while the product that ships is one none of them saw. This paper sets out one integrated control framework for banks and insurers. One AI product registry records every model, agent, prompt, tool, data product and vendor in use; if it is not in the registry, it does not run in production. One risk assessment carries the risk, legal and cyber questions together, sets a tier and gates every release, with independent second-line approval for the higher tiers. Guardrails, identity and monitoring enforce at runtime what the registry approved, and every decision leaves evidence. The three lines of defense run through every control point: the first line registers and operates, the second line sets standards and approves independently, internal audit tests the gate and the evidence, and agents never approve.

WHO IT'S FOR	Chief risk and compliance officers, general counsel and regulatory lawyers, CISOs, chief audit executives, and the CIOs, CDAOs and product owners who build AI at banks and insurers.
YOU'LL BE ABLE TO	Stand up one AI product registry and tiering questionnaire, run a production gate with evidence by tier and named independent approvers, assign the three lines at every control point, map US, EU, UK, Canadian and Singapore obligations and the OWASP and MITRE ATLAS threats to controls, and report to the board from the same record.
SHORT ON TIME?	Read the executive summary and sections 03 and 06, then the 90-day plan in section 17 and the appendix for your role.

The full paper follows.

EXECUTIVE SUMMARY

One registry, one gate, one evidence trail.

AI governance fails when risk, legal and security each run their own process. Three reviews ask three sets of questions about three different descriptions of the same system, on three calendars. Each can pass while the product that reaches production is one none of them saw. This paper sets out one integrated control framework for banks and insurers, written from three seats at once: the chief risk and compliance officer, the technology and regulatory lawyer, and the chief information security officer.

The argument in four sentences

- **Register every AI product once.** One AI product registry holds the record for each model, agent, prompt, tool, data product and vendor in use, with a named owner. If it is not in the registry, it does not run in production.
- **Gate production with one risk assessment.** A single assessment carries the risk, legal and cyber questions together. It sets a tier, and the tier decides which evidence is needed and who must approve, independently of the team that built it.
- **Enforce at runtime what the registry approved.** Guardrails, identity, data controls and monitoring read the approved record. A release ID that is not approved gets no tools and no data.
- **Leave evidence at every decision.** Intake, tiering, approval, runtime, incident, change and retirement each write to one evidence trail that a supervisor, an auditor or a court can follow.

The three lines, integrated rather than stacked

The first line (business and technology owners) registers, assesses, operates the controls and monitors. The second line (risk, model risk, compliance, legal and the CISO office) sets the standards, confirms the tier, challenges the evidence and approves the higher tiers independently; it never owns delivery. The third line (internal audit) tests the gate and the evidence afterwards. All three work on the same record. Agents and automation can prepare evidence and block a release. They never approve one.

FROM PRACTICE

In my experience, the organizations that governed AI well were not the ones with the longest policies. They had a product registry, and nothing reached production without a completed risk assessment attached to the registry entry. That one rule did more than any committee, because it put the risk, legal and security questions in front of the people building the product, before release, in one place.

NOT LEGAL ADVICE

This paper describes a control framework and how to evidence it. It is not legal advice, a legal opinion or a conformity assessment. Laws and supervisory expectations change; each regulatory statement carries a dated primary source in the register at the end, checked through 20 November 2026. Your counsel decides what applies to you.

Ten decisions for the executive committee.

The framework is small enough to decide in one meeting. These are the decisions that make it real, with the section that explains each.

DECISION	WHAT YOU ARE AGREEING TO	SECTION
1 One registry	Every AI product and component in production has one registry record and a named first-line owner	04
2 One assessment	Risk, legal and cyber questions sit in one questionnaire that sets a tier from 1 (critical) to 4 (low)	05
3 No record, no production	Deployment and the runtime gateway refuse anything without an approved registry release	06
4 Independent approval	Tier 1 and Tier 2 releases need second-line approval; the approver did not build it and no agent approves	03, 06
5 Exceptions expire	Every exception has an owner, a compensating control and an end date, and the second line accepts it	06
6 Controls outside the model	Identity, tool limits and data rules are enforced by systems the model cannot talk its way past	08, 09
7 Tested fallback	No fallback model or provider serves a regulated use until it has passed the same evaluation	10
8 One incident clock table	AI incidents use the existing incident process, with every regulatory and contractual clock mapped	13
9 Committee reporting	The board's risk and audit committees get one AI report from the registry each quarter	03, 14
10 Change re-enters the gate	A model, provider, prompt, tool or data change that is material goes back through the gate	15

Where to start, by role

IF YOU ARE	READ FIRST
Chief risk or compliance officer	Sections 03, 05, 06 and 14, then Appendices B and C
General counsel or regulatory lawyer	Sections 11 and 12, then the source register
CISO	Sections 07 to 10 and 13, then Appendices D and E
CIO, CDAO or product owner	Sections 02, 04, 06 and 16, then the 90-day plan in section 17
Chief audit executive	Sections 03 and 14, then Appendix C

For the supervisory evidence behind each control, this paper relies on AI Governance Evidence, which covers what US bank and insurance reviewers inspect in detail. This paper is the operating framework that produces that evidence.

Three reviews of three different systems.

Most firms did not design their AI governance. They inherited it. Model risk extended its validation process, legal added an AI question to the new-product checklist, and security added AI to the vendor and architecture reviews. Each process is reasonable alone. Together they fail in predictable ways.

FAILURE	WHAT IT LOOKS LIKE	WHY IT HAPPENS
Different units of review	Model risk validates the model; legal reviews the use case; security reviews the vendor. Nobody reviews the deployed product: model, prompt, retrieval, tools, permissions and workflow together	Each function reviews the object it already knows
Different inventories	Three lists of AI that do not reconcile, none complete, none linked to what is deployed	Each review keeps its own register
Sequential calendars	A product waits weeks for the third review, then ships a different version from the one first reviewed	Reviews are queues, not a gate on a version
Approval without enforcement	The approval says read-only; the deployed agent holds a write tool	Nothing at runtime reads the approval
Ownerless change	A provider updates its model, or a prompt is edited, and nobody re-assesses	Change is defined for code, not for models, prompts or data
Shadow AI	Teams use AI features in tools already bought, outside every review	Intake only catches new purchases
Evidence after the fact	Before an examination, a team spends weeks rebuilding who approved what	Decisions were made in email and meetings

Supervisors are now asking about exactly these seams. In April 2026 the US banking agencies issued SR 26-2, revised model risk guidance that leaves generative and agentic AI outside its scope and points banks to their broader risk management and governance for them.[1] In November 2026 the OCC gave its examiners a framework for reviewing banks' use of AI, organized around purpose, risk, control, oversight, resilience and assurance.[2] The question an examiner asks is no longer whether you have an AI policy. It is whether you can show the controls work on the thing you deployed.

AI changes three things existing processes assume.

Risk, legal and security processes were built for systems that behave the same way every time, change only when someone releases code, and act only through screens people use. Generative and agentic AI breaks each assumption.

ASSUMPTION	WHAT CHANGED	CONSEQUENCE FOR GOVERNANCE
The system is deterministic	The same input can produce different outputs, and quality depends on context the model is given at run time	Validation once is not enough; evaluation must be repeated and monitoring must watch outputs
Change means a code release	Behavior changes with a provider's model update, a prompt edit, a new document in the index or a new tool	The registry must version every component, and change control must cover all of them
People take the actions	Agents call tools, write records and send messages on a person's behalf	Authority, identity and approval must be enforced per action, outside the model

Why the three lenses must be joined

Each AI failure is usually all three kinds of failure at once. An assistant that discloses a customer's account details to the wrong caller is a security incident, a privacy breach that may start a notification clock, and an operational risk event. A credit model that relies on an unapproved data field is a fair lending question, a model risk finding and a data governance gap. If three teams own three slices of the response, the firm learns slowly and reports inconsistently.

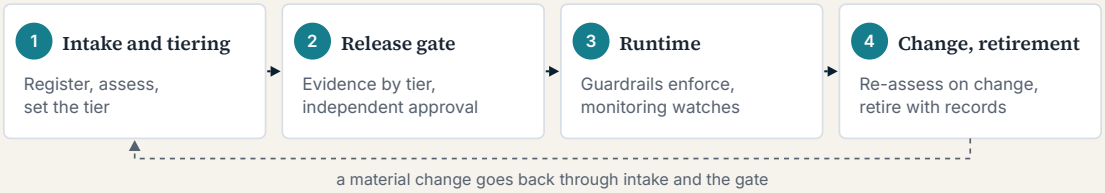
The answer is not a new AI department. It is one record, one assessment and one evidence trail that each function contributes to and relies on, with clear rules about who decides what. The rest of this paper builds that, control point by control point.

> **Governance is the set of decisions that can stop a release, and the evidence that they were made.**

SCOPE Banking and insurance, with a short healthcare note in section 11. Jurisdictions: US federal and state, the EU, the UK, and short notes on Canada and Singapore. Uses: any AI that informs or takes a business decision, including assistants, agents, document extraction, scoring models and AI features inside purchased software.

Three lenses, four control points, three lines.

The framework has three parts. Three lenses decide what is asked. Four control points decide when it is asked. Three lines decide who asks, who answers and who checks.



FIRST LINE	Registers the product, answers the questions	Builds the evidence pack for its tier	Operates guardrails, monitors, raises incidents	Proposes changes, retires and archives
SECOND LINE	Owns the tier standard, confirms Tier 1 and 2	Challenges evidence, approves Tier 1 and 2	Reviews signals against appetite, escalates	Approves material change, exceptions
THIRD LINE	Tests that every product in use is registered	Re-performs the gate on a sample	Tests monitoring and incident handling	Tests change records and retirement

ONE REGISTRY RECORD · ONE RISK ASSESSMENT · ONE EVIDENCE TRAIL

Exhibit 2.1: Four control points, three lines, one record. The second line owns the standard and approves the higher tiers; it never owns delivery.

LENS	ASKS	OWNED IN THE SECOND LINE BY
Risk	What can go wrong for customers, the firm and its counterparties, how likely and how bad, and is it within appetite	Chief risk officer, model risk management, operational risk
Legal	Which laws, rules, contracts and rights apply, what they require, and who carries liability	General counsel, chief compliance officer, privacy officer
Cyber	How it can be attacked or misused, what it can reach, and how we would know	Chief information security officer and the security architecture team

AI risk is not a new risk category. The OCC tells its examiners not to treat it as one, and to expect findings in existing areas such as operational, compliance or third-party risk.[2] Map each AI risk on the registry record to the firm's existing risk taxonomy and appetite statements, so AI exposure rolls up into the same reports as everything else.

What each control point decides.

Each control point is a decision with an owner, an input and an output written to the registry. The lenses are questions inside each decision, not separate reviews.

CONTROL POINT	DECISION	RISK LENS	LEGAL LENS	CYBER LENS	EVIDENCE WRITTEN
1 Intake and tiering	Register it, and how much control it needs	Impact, autonomy, affected people	Legal triggers, jurisdictions, data rights	Attack surface, data sensitivity	Registry record, questionnaire, tier
2 Release gate	May this version run in production, for this scope	Evaluation, validation, fallback	Sign-off, notices, impact assessments	Red team, access tests, vendor review	Release decision with approvers
3 Runtime	Is it behaving as approved	Drift, override and error rates	Complaints, disclosures, retention	Attacks, leakage, cost spikes	Monitoring record, incidents
4 Change and retirement	Does a change need a new gate; is it safe to stop	Material change test	Records, legal hold, contracts	Credentials, data removal	Change and decommission records

Three design rules

- **The product is the unit of approval.** Approve the deployed composition: model and version, prompts, retrieval sources, tools, permissions, data products, user interface and the workflow around it. A model approved in isolation is not an approved product.
- **The tier sets the depth, not the route.** Every product passes the same four control points. A Tier 4 internal drafting aid passes them in a day with first-line evidence. A Tier 1 agent that moves money needs independent validation and committee visibility. Same route, different depth.
- **Runtime reads the registry.** The release ID approved at the gate is what the gateway, identity and monitoring check on every call. An approval that runtime cannot read is a document, not a control.

This is consistent with the NIST AI Risk Management Framework, whose Govern function runs through Map, Measure and Manage rather than standing beside them.[24] The registry is where those four functions meet one product.

RELATION TO OTHER RJL PAPERS

The Four-Part System and The Reference Implementation describe the architecture this runs on. Agent Onboarding describes the agent contract the registry stores. The Production-Readiness Diagnostic is a ten-question check of whether this framework is working.

The three lines, defined for AI.

Banks and insurers already organize risk in three lines; the Institute of Internal Auditors now calls it the Three Lines Model.[33] AI does not need a new model. It needs the existing one applied precisely, so that the people who build AI are not the only people judging whether it is safe to rely on.

LINE	WHO	DOES	NEVER DOES
First line	Business owner of the use; product and technology owners who build and run it; data and model owners	Registers the product, completes the risk assessment, builds the evidence, operates the controls, monitors, raises incidents, proposes changes, retires	Approves its own Tier 1 or Tier 2 release; accepts its own critical exception
Second line	Risk and model risk management, compliance, legal and privacy, the CISO office	Sets the standards and the tiering method, confirms tiers, challenges evidence, approves Tier 1 and 2 releases independently, sets risk appetite, accepts exceptions, reports to committees	Builds or runs the product; owns delivery dates; writes the evidence it then approves
Third line	Internal audit	Gives independent assurance: tests that the registry is complete, re-performs the gate on samples, tests runtime controls and evidence, reports to the audit committee	Designs or approves controls it will later audit

Where the CISO and legal sit

Firms differ. In many, the CISO office runs security operations, which is first-line work, and sets security standards, which is second-line work. Name which role does which. The person who approves a Tier 1 release from a security standpoint should not be the person who built its controls. The same applies to legal: product counsel advising the build team is first-line support; the legal or compliance sign-off at the gate must be independent of that advice.

Model risk management keeps its role for models in scope of SR 26-2. Because generative and agentic AI are outside that guidance, name explicitly which second-line function validates them.[1] AI Governance Evidence calls this the gap to close first.

RULE One person can hold roles in more than one line in a small firm. One person cannot both build and approve the same release, and one decision cannot be both made and tested by the same line.

Who does what at each control point.

The table below is the summary RACI for the seven decisions this paper covers. Each later section repeats the rows for its control point in more detail. Accountability stays in the first line for the use and its outcome; the second line holds independent approval for the higher tiers.

ACTIVITY	BUSINESS OWNER (1L)	PRODUCT AND TECH OWNER (1L)	RISK AND MODEL RISK (2L)	LEGAL AND COMPLIANCE (2L)	CISO OFFICE (2L)	INTERNAL AUDIT (3L)
Intake: register the product	**A**	R	C	C	C	I
Tiering: set and confirm the tier	**A**	R	**Ap**	C	C	I
Release gate: Tier 3 and 4	**A**	R	I	C	C	T
Release gate: Tier 1 and 2	**A**	R	**Ap**	**Ap**	**Ap**	T
Runtime: operate and monitor	**A**	R	C	I	C	T
Incident: declare and contain	**A**	R	C	C	R	I
Incident: regulatory notification	C	C	C	**A**	R	I
Change: material change decision	**A**	R	**Ap**	C	C	T
Retirement: decommission with records	**A**	R	I	C	C	T

R does the work · A accountable for the outcome · Ap approves independently · C consulted · I informed · T tests afterwards. One A per row. 1L, 2L, 3L: first, second and third line.

How the lines work on one record

In the integrated model the three lines do not each run a review. They each work on the same registry record. The first line fills it in. The second line reads the same record, challenges it in comments that stay attached, and approves or declines with a reason. Internal audit samples the same records later and re-performs the decision. Nobody keeps a private copy of the facts.

- **Challenge is recorded.** A second-line question and the first-line answer stay on the record. An approval with no recorded challenge on a Tier 1 product is itself a finding.
- **Approval is bound to a version.** The approver approves release ID, scope and expiry, not the product in general. A changed component voids the approval.
- **Audit tests the system, not a sample of slides.** Because every decision is in the registry, audit can select its own sample from the full population.

Independence, two people, and where agents fit.

Independence is the property that makes a review worth having. Write the rules down, and have the registry enforce the ones a system can check.

RULE	WHAT THE REGISTRY CHECKS
Two people for every release	Requester and approver are different named people; a Tier 1 release has at least three distinct approvers across first and second line
No self-approval	The approver is not the builder, the requester, or in the builder's reporting line for Tier 1 and 2
Second line independent of delivery	A second-line approver holds no delivery objective for the product and is not paid on its launch
Exceptions accepted above the requester	A critical exception is accepted by the second line, never by the team that needs it
Audit independent of design	Internal audit did not design or approve the controls it tests
Agents never approve	No service account, agent identity or automated job can hold an approval role; approvals require a signed-in person

What agents and automation may do

Automation is welcome in governance, inside a boundary. It may assemble evidence, run evaluations and red-team suites, compare the deployed composition with the approved one, check that required fields and signatures exist, draft a summary for an approver, and block a release or a call that fails a rule. It may not approve, accept a risk, accept an exception, close an incident, or count as one of the two people. An AI reviewer that scores an answer is an evaluation tool; its output is evidence a person weighs, not a decision.

This mirrors where supervisors are heading. The Financial Stability Board's final sound practices of October 2026 call for clear triggers for when an agent must obtain human approval, and tested ways to stop it.[7]

BOARD AND COMMITTEE REPORTING

The registry produces one quarterly AI report for the board's risk committee and audit committee: inventory by tier, releases approved and declined, open exceptions by age, incidents, material changes and audit findings. Section 14 sets out its contents and the metrics behind it.

Register the product and every part it is made of.

The registry is the system of record for AI in production. It is not a spreadsheet maintained for an examination. It is the list that deployment and runtime read, so that an unregistered product cannot run.

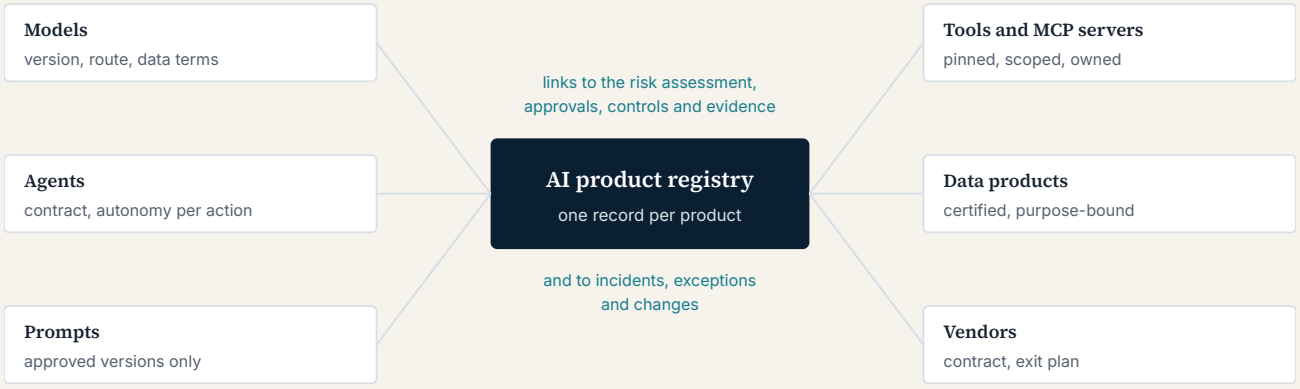


Exhibit 4.1: The product is the unit of record. Its components are registered too, and each component change is a change to every product that uses it.

COMPONENT	WHY IT IS REGISTERED ON ITS OWN	MINIMUM IT CARRIES
Models	One model may serve many products; a provider update changes all of them	Provider, version or pin, route, data processing terms, approved uses, evaluation record
Agents	An agent acts; its authority must be bounded per action	Contract, owner, tools, autonomy level per action, limits, kill switch
Prompts	A prompt edit can change behavior as much as a model change	Version, author, approver, evaluation run on promotion
Tools and MCP servers	A tool is a capability the model can invoke; a changed tool is a changed product	Owner, scopes, pinned digest, side effects, approval rule
Data products	Retrieval and features decide what the model knows	Owner, classification, permitted purposes, certification state
Vendors	Accountability stays with you when a vendor supplies the AI	Contract terms, sub-processors, incident duty, exit plan, fallback

The minimum record, and the rule that makes it matter.

Appendix A gives the full template field by field. These are the fields every record needs before a product can pass the gate, grouped by the lens that relies on them.

GROUP	FIELDS	USED BY
Identity	Registry ID, product name, release ID, environment, components with versions, first-line owner, technical owner	Everyone; the release ID is what runtime checks
Purpose and scope	Intended use, out-of-scope uses, users, affected people, decisions informed, autonomy level per action, jurisdictions	Risk, legal
Risk assessment	Questionnaire answers, tier, tier rationale, second-line confirmation, risk appetite reference	Risk
Legal	Applicable obligations from the control map, required notices, impact assessments, contract references, retention class, legal hold flag	Legal and compliance
Security	Threat model, data classes reached, identities and scopes, tools and side effects, red-team result, exposure	CISO office
Release	Evidence by tier, approvers and dates, conditions, expiry, exceptions with end dates	All three lines
Operation	Monitoring signals and thresholds, incidents, changes, fallback, next review date, retirement plan	First line, then all

FROM PRACTICE

In my experience, the registry worked when it was the only door to production. Teams stopped arguing about governance once the deployment pipeline simply asked for a registry ID with an approved risk assessment, and refused without one. The assessment became part of building the product, not a review after it. The registries that failed were the ones kept beside delivery, filled in afterwards, and trusted by nobody.

No record, no production

Make the rule mechanical. The deployment pipeline refuses to promote a release without an approved registry release ID. The runtime gateway refuses tool and data calls from a release ID that is not approved, expired or revoked. Section 06 covers the gate itself. Section 08 covers the runtime check.

Find the AI you did not register.

A registry that lists only what teams declared will be incomplete on day one. Shadow AI is any AI in use that the registry does not know about: an AI feature switched on in software you already license, a team calling a model API with a corporate card, a browser extension, a script that sends documents to a public assistant.

DISCOVERY SOURCE	WHAT IT FINDS	OWNER
Network and proxy logs	Calls to known model and AI service endpoints from corporate devices and workloads	CISO office
Cloud and API key inventory	Model service accounts, keys and spend in cloud accounts	Technology owner, CISO office
Procurement and expense data	Purchases and renewals of AI tools, and AI features added to existing contracts	Procurement, second-line risk
SaaS administration consoles	AI features enabled in collaboration, CRM and productivity suites	Technology owner
Code and pipeline scans	Model SDKs, prompts and agent frameworks in repositories	Product and tech owner
Data security posture tools	Sensitive data copied into AI services or vector stores	CISO office, data owner
Endpoint and developer tools	Local MCP servers, coding agents and browser extensions that call models	CISO office
Attestation	A short quarterly attestation by business heads of AI in use in their area	Business owner

Treat each finding as an intake, not a violation. Register it, tier it, and either bring it through the gate or switch it off on a date. Report the count of unregistered AI found and closed to the committee: a falling number is evidence the registry is becoming the door it should be.

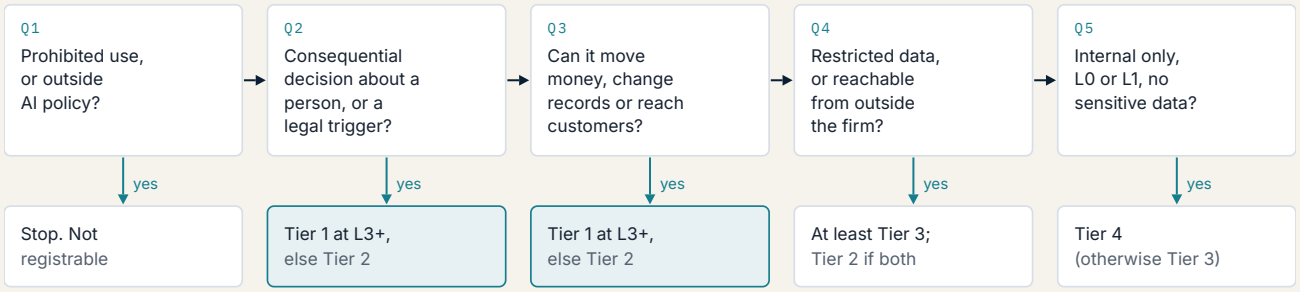
Supervisors expect an inventory. Colorado’s insurance regulation requires one for external data, algorithms and models, and Singapore’s new guidelines expect financial institutions to identify AI use and keep inventories.[19, 45]

THIRD-LINE TEST Internal audit compares discovery results against the registry each year and reports the unregistered rate. The test is whether the population is complete, not whether the records that exist are tidy.

Six questions decide how much control a product needs.

Intake turns a request into a registry record. Tiering decides how deep each later control goes. Use one questionnaire that risk, legal and security all wrote, so the first line answers once and every lens reads the same answers.

DIMENSION	WHAT IS ASKED	RAISES THE TIER WHEN
Use and impact	Which decision or task it informs, how consequential an error is, how reversible	It informs credit, pricing, claims, employment, fraud holds or access to a service
Autonomy	What it may do per action on the L0 to L4 scale, and who reviews	It acts without per-action approval (L3 or L4) on money, records or customers
Data sensitivity	Which data classes it reads, writes or sends to a provider	It touches nonpublic personal information, health data, credentials or regulated records
Affected people	Who is affected, how many, whether they are customers, employees or the public, and whether they know	Consumers are affected, at scale, or without a human contact
Legal triggers	Which named obligations apply: fair lending, unfair discrimination, AI Act high-risk, automated decision rules, sector law	Any named trigger applies; counsel decides
Attack surface	Who can reach it, which tools and outbound channels it holds, which untrusted content it reads	It is reachable from outside the firm, reads untrusted content and holds tools



A no moves to the next question. Ask every question; the highest tier any answer produces wins.
The second line confirms every Tier 1 and Tier 2 result. L0 to L4 is the autonomy scale from Agent Patterns.

Exhibit 5.1: The tiering flow. Appendix B holds the full questionnaire and scoring.

The tiers follow the four illustrative tiers in the OCC's November 2026 examiner framework, from Tier 1 critical to Tier 4 low, which suggests classifying by consequence, transfer of agency and degree of human oversight rather than by technical sophistication.[2] Agent Patterns maps its autonomy levels to the same tiers.

Tiering is a decision with an owner.

The first line proposes the tier from its answers. For Tier 1 and Tier 2 the second line confirms it before evidence work starts, because a tier set too low is the cheapest way to skip a control.

TIER	PROFILE	EXAMPLE IN BANKING OR INSURANCE
Tier 1 Critical	Acts directly with little reversibility and high customer or financial impact, or decides about people with no per-decision human review	An agent that releases payment holds; automated claim denial; real-time credit decisioning
Tier 2 High	High-weight decision support or direct action with a person on the loop; or any named legal trigger	Underwriting assistance; AML alert disposition support; customer-facing service agent that can change records with approval
Tier 3 Moderate	Decision support with a person in the loop, reversible, limited data sensitivity	Claims file summarization for an adjuster; credit memo drafting
Tier 4 Low	A person fully in control, no customer or sensitive data	Code assistance; internal policy search; meeting notes

Several of these examples follow those the OCC framework gives for its illustrative tiers.[2]

ACTIVITY	BUSINESS OWNER (1L)	PRODUCT AND TECH OWNER (1L)	RISK AND MODEL RISK (2L)	LEGAL AND COMPLIANCE (2L)	CISO OFFICE (2L)	INTERNAL AUDIT (3L)
Submit intake and answer the questionnaire	**A**	R	I	I	I	I
Identify legal triggers and jurisdictions	**A**	R	C	R	I	I
Describe attack surface and data reached	**A**	R	I	C	C	I
Propose the tier	**A**	R	C	C	C	I
Confirm Tier 1 and Tier 2	I	I	**A**	C	C	T
Maintain the tiering standard	I	I	**A**	R	R	T

R does the work · A accountable for the outcome · Ap approves independently · C consulted · I informed · T tests afterwards. One A per row. 1L, 2L, 3L: first, second and third line.

Three tiering mistakes

- **Tiering the model, not the use.** The same model can sit in a Tier 4 drafting aid and a Tier 1 claims agent. Tier the product.
- **Treating "a human reviews it" as the answer.** A reviewer with no time, no information or no authority to say no is not a control. Record the review rate and the override rate, and test them at runtime.
- **Never re-tiering.** A tool added, a new user group or a new jurisdiction can move a product up a tier. Section 15 makes that a change trigger.

Evidence by tier.

The gate is the decision that a specific release may run in production for a stated scope, until a stated date. The tier sets the evidence. The evidence is assembled once, on the registry record, and each lens reads the part it needs.

EVIDENCE	TIER 1	TIER 2	TIER 3	TIER 4
Workflow evaluation against criteria set before testing	Independent	Required	Required	Smoke test
Critical-failure veto (no average offsets a critical failure)	Required	Required	Required	Required
Red team against the threat model (Appendix D)	Independent	Required	Gateway cases	Not required
Fairness and unfair-discrimination testing where a person is affected	Required	Required	If a trigger applies	Not required
Legal and compliance sign-off on the control map	Required	Required	Checklist	Checklist
Privacy impact assessment, or DPIA where GDPR applies	Required	Required	Screening	Screening
Vendor and model provider review	Full	Full	Standard	Standard
Access, revocation and tool-limit tests on the deployed route	Required	Required	Required	Required
Human oversight exercise in the real interface	Required	Required	If L1 or L2	Not required
Monitoring plan with thresholds and owners	Required	Required	Required	Basic
Rollback and kill-switch test; tested fallback	Required	Required	Rollback	Rollback

“Independent” means performed or re-performed by people outside the build team: model risk, a separate validation group or a qualified third party. For banks, SR 26-2’s independent validation expectations apply to models within its scope; for generative and agentic capabilities outside it, the gate should name who plays that role.[1]

DESIGN AND OPERATING EVIDENCE

Every control needs both: the design (policy, configuration, contract) and proof that it operated on the deployed route (allowed and denied requests, a failed test that was fixed). AI Governance Evidence describes what reviewers accept for each.

Who approves, and what they approve.

An approval is a signed decision on a release ID, a scope and an expiry. It names the conditions that keep it valid. It is never an approval of a product in general.

TIER	APPROVERS (ALL NAMED PEOPLE)	VISIBILITY
Tier 1	Business owner and product owner (first line); risk or model risk after independent validation; legal or compliance; CISO office. Three or more distinct people	Reported to the board's risk committee before or at launch
Tier 2	Business owner and product owner; one second-line approver each from risk, legal or compliance, and the CISO office	Quarterly committee report
Tier 3	Business owner and product owner, with a second-line reviewer consulted on any open legal or security question	Registry; sampled by second line
Tier 4	Two first-line people: the owner and someone other than the builder	Registry; sampled by second line

ACTIVITY	BUSINESS OWNER (1L)	PRODUCT AND TECH OWNER (1L)	RISK AND MODEL RISK (2L)	LEGAL AND COMPLIANCE (2L)	CISO OFFICE (2L)	INTERNAL AUDIT (3L)
Assemble the evidence pack	**A**	R	I	I	I	I
Run independent validation (Tier 1)	I	C	**A**	I	C	T
Legal and compliance sign-off	I	C	C	**A**	I	T
Security sign-off and red-team result	I	C	I	I	**A**	T
Approve Tier 1 or Tier 2 release	**A**	R	**Ap**	**Ap**	**Ap**	T
Approve Tier 3 or Tier 4 release	**A**	R	I	C	C	T
Bind the approval to the release ID in the registry	I	R	**A**	I	C	T

R does the work · A accountable for the outcome · Ap approves independently · C consulted · I informed · T tests afterwards. One A per row. 1L, 2L, 3L: first, second and third line.

An approval can be conditional: approved for a smaller user group, a lower autonomy level or read-only use while a gap closes. Conditions are fields on the record, and runtime enforces them like any other scope limit.

The release approval also states the operating conditions: the user population, allowed actions and autonomy level per action, data boundary, monitoring thresholds that trigger review, and the next review date. Funding a prototype, allowing production reliance and accepting operational handover are three different decisions; the gate is only the second.

Exceptions expire. Records come first.

A gate that never says no is a calendar. A gate that never grants an exception is ignored. The answer is a narrow, visible exception process with an end date.

ELEMENT	RULE
What can be excepted	A missing or failed control where a compensating control holds the residual risk within appetite. Never access control, isolation, action authority, integrity or required legal notices
Who requests	The first-line owner, with the compensating control and the plan to close
Who accepts	The second line, above the requester; for Tier 1, the chief risk officer or a delegate named in policy
How long	Suggested: Tier 1 up to 30 days, Tier 2 up to 90 days, Tier 3 and 4 up to 180 days, with one renewal and a new acceptance. Set your own limits in policy; no regulation fixes them
What happens at expiry	The registry marks the release non-compliant; the gateway restricts or stops it unless the exception was closed or renewed
Who sees it	Open exceptions by tier and age go to the quarterly committee report; audit tests closure

FROM PRACTICE

In my experience, the risk assessment worked as a gate because it was attached to the product record, not to a meeting. The release went ahead when the record was complete and the right people had signed it, and not before. When a team needed an exception, it was written on the same record with an end date, so everyone could see what was being carried and for how long.

No record, no production, enforced

- The deployment pipeline checks for an approved registry release ID before it promotes a build.
- The runtime gateway checks the release ID on every tool and data call, and denies an unknown, expired or revoked release.
- A changed component (model version, prompt, tool digest, data product version) produces a new release ID that is not approved until it passes the gate.
- Every refusal is logged against the registry ID, so the committee sees attempts as well as approvals.

What attackers and accidents do to AI products.

Threat modeling for AI starts from the same question as for any system: what can reach it, what can it reach, and what happens if it is wrong or turned against you. The table maps the threats to the OWASP Top 10 for LLM Applications, 2026 edition, the OWASP Top 10 for Agentic Applications, and MITRE ATLAS.[28, 29, 30]

THREAT	HOW IT HAPPENS	OWASP 2026	PRIMARY CONTROLS
Direct prompt injection	A user writes instructions that override the system's rules	LLM01	Input screening; authority outside the model; output checks
Indirect prompt injection	Instructions hidden in a document, web page, email or tool result the model reads	LLM01; ASI01 goal hijack	Treat retrieved content as data; result screening; tool gating
Sensitive data leakage	The model reveals data the user may not see, or sends it to a provider or tool	LLM02	Retrieval filtered by the user's rights; masking; route approval by data class
Excessive agency	The agent holds more tools, scopes or autonomy than the task needs	LLM03	Least privilege tools; autonomy per action; approval binding
Supply chain	A compromised model, package, plugin or MCP server	LLM04	Registry pinning; provenance; tool digests; vendor review
Data and model poisoning	Tampered training data, fine-tuning sets or retrieval corpora	LLM05	Source control on corpora; change review; evaluation on change
Denial of wallet	Loops, floods or large contexts that run up cost or exhaust quota	LLM06	Run limits; spend caps; rate limits; alerts
Misinformation	Confident, wrong answers people rely on	LLM07	Grounding checks; citations; incomplete state; human review
Hidden context exposure	System prompts or hidden context disclosed, including secrets placed in them	LLM08	No secrets in prompts; output screening
Vector store weaknesses	Embeddings that bypass access rules or leak across tenants	LLM09	Filter before ranking; tenant isolation
Insecure output handling	Model output executed or rendered without validation	LLM10	Typed tools; no free-text SQL or code execution; encoding

Agents add identity, tools and time.

An assistant that only answers can leak and mislead. An agent that acts can also do damage at machine speed, under someone else's identity, while nobody is watching. The OWASP agentic list adds threats that a chat-only threat model misses.[29]

AGENTIC THREAT	EXAMPLE	CONTROL
Goal hijack through content	A supplier invoice tells the payment agent to change the bank account on file	Read and write separated; write tools need approval bound to the exact change
Tool misuse	A summarization agent uses an email tool to send a file outside the firm	No outbound channel unless the contract names it; allowlisted recipients
Identity and privilege abuse	An agent acts with a service account broader than the person it serves	Delegated, narrowing identity per request (section 09)
Memory poisoning	A planted instruction saved in shared memory affects later sessions	Memory with declared readers and writers, provenance and expiry
Cascading failure	One agent's bad output becomes another agent's trusted input	Each hop re-checks authority; limits per chain
Rogue or unregistered agents	An agent deployed outside the registry calls production tools	Gateway refuses unknown release IDs and identities

Using ATLAS and NIST in the threat model

MITRE ATLAS catalogs adversary tactics and techniques against AI systems in the style of ATT&CK, including techniques aimed at agents and their tools, with case studies. Use it to write red-team scenarios and to map detections in the security operations center.[30] NIST AI 100-2 gives the vocabulary for attacks on predictive and generative systems, and NIST's draft Cyber AI Profile maps AI to the Cybersecurity Framework; treat that profile as a draft.[26, 27]

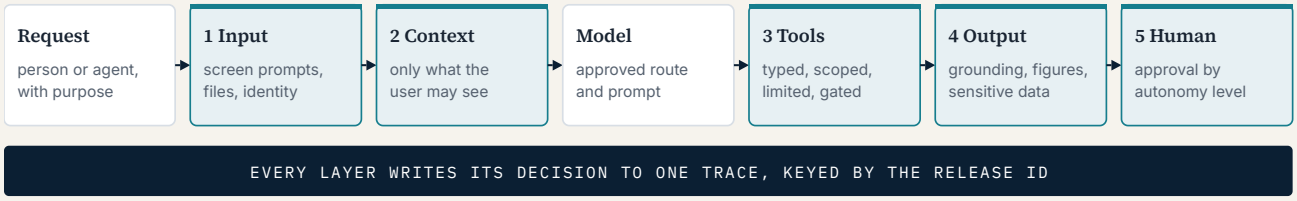
Supervisors are watching the attacker side too. In May 2026 NYDFS warned regulated entities that frontier AI models speed up finding vulnerabilities and writing exploits, and asked them to shorten remediation timelines, coordinate with critical third parties and add human review of AI-generated code before production. The letter imposes no new requirements.[23]

OWNER The CISO office owns the threat model standard. The first line writes the threat model for each Tier 1 to 3 product on the registry record; the CISO office challenges it and signs the security part of the gate.

08 / GUARDRAILS, LAYER BY LAYER

Enforce at runtime what the registry approved.

A guardrail is a control that runs on every request. Put it where the model cannot negotiate with it. Each layer stops some failures and misses others; that is why there are five.



Layers 1 to 3 are enforced outside the model. The model is never the control; it is what is controlled.

Exhibit 8.1: Guardrail layers. The table below says what each stops, and what it does not.

Run layers 1 to 3 in one gateway in front of models, tools and data, not separately inside each application. One implementation means one set of rules to test, one log format and one place for the release ID check.

LAYER	STOPS	DOES NOT STOP
1 Input screening	Known injection patterns, oversized or malformed input, files of the wrong type, requests from unregistered releases	Novel injections; instructions hidden in content retrieved later
2 Context and retrieval	Data the user may not see; stale or unapproved sources; cross-tenant passages; restricted labels reaching an unapproved route	A wrong answer built from permitted data
3 Tool and action gating	Tools outside the contract; arguments outside type and range; free-text SQL; writes without bound approval; loops past run and spend limits	A permitted action taken for a bad reason inside its limits
4 Output checks	Unsupported claims, figures that do not match the governed metric, sensitive data in the answer, unsafe rendering	Subtle errors that look grounded
5 Human approval	Consequential actions a reviewer would refuse	Anything the reviewer has no time, information or authority to judge

Human approval follows autonomy, per action.

Human oversight is a guardrail with the highest cost and the weakest consistency. Use it where it adds judgment, and make it real where you use it.

AUTONOMY LEVEL	THE PERSON	GUARDRAIL REQUIREMENT
L0 Assist	Decides and acts on a suggestion	Output checks; clear labels for generated content; usage logging
L1 Approve each	Approves every action before it runs	Approval bound to the exact proposal and parameters, single use, with expiry
L2 Approve the consequential	Approves money, customer and record changes	Consequence defined per action in the contract; reviewer has the information to say no
L3 Monitor	Reviews samples, exceptions and metrics	Sampling plan, alerting, fast rollback, kill switch
L4 Bounded autonomy	Owns the outcome, reviews performance	Hard limits on spend, scope and data; periodic re-approval; tested stop

Grounding, figures and sensitive data

- **Grounding.** Each factual statement should be traceable to a permitted source version. Where evidence is missing, stale or conflicting, the product returns a defined incomplete state rather than a guess.
- **Figures.** Numbers that matter come from governed metrics computed outside the model, carried with their version. A reportable number the model produced itself fails the check.
- **Sensitive data.** Output screening for personal, health and account data runs after generation as well as before retrieval, because models can recombine permitted fragments into a disclosure.

Guardrails are evidence as well as controls. Every block, mask, refusal and approval is written to the trace with the rule and version that made it. Monitoring in section 13 reads those records; audit samples them.

Feed security operations

Send guardrail events to the security operations center as well as the product team: blocked injections, tool denials, unknown release IDs, revocations and unusual data volume. Write detections for them, mapped to MITRE ATLAS techniques where one fits, and include AI products in the existing incident on-call rota.[30]

WHAT NOT TO RELY ON — A system prompt that tells the model to obey rules is guidance to the model, not a control. Test every guardrail as if the model will try to get around it, because an attacker will try through the model.

09 / IDENTITY, ACCESS AND DATA PROTECTION

Every call carries a person, an agent and a purpose.

Most AI data incidents are access control failures in a new place. The fix is to carry identity and purpose all the way to the data, and to enforce there.

CONTROL	WHAT GOOD LOOKS LIKE
Agent identity	Every agent has its own registered identity, separate from any person and from shared service accounts, tied to its registry record
Delegation	A request carries both the person and the agent, with audience and purpose, in a token that can only narrow as it passes agent to agent; RFC 8693 token exchange is one standard shape for this[32]
Purpose binding	The purpose declared at intake is checked on every call; a claims purpose cannot read marketing data
Least privilege tools	Tools are typed, scoped to the minimum operation and listed in the agent contract; no general-purpose shell, browser or SQL tool in a regulated workflow without its own gate
Row and field controls	Enforced in the data layer or the gateway in front of it, with the same rules for direct queries, caches and vector indexes
Masking and tokenization	Restricted identifiers are masked or tokenized before they reach a model route not approved for them; detokenization needs two people
Secrets	Credentials live in a secrets manager, never in prompts, code, logs or the browser; agents receive short-lived scoped tokens, not keys
Revocation	A revoked person or agent loses access within a stated window, and the window is tested
Network egress	Agent runtimes reach only the gateway and approved endpoints; no general internet access unless the contract names it

For Model Context Protocol servers, follow the protocol's security best practices: authorization on every call, no token passthrough to downstream services, and protection against the confused-deputy problem where a server acts with its own broader rights.[31]

REGULATORY ANCHOR

NYDFS's October 2024 letter places AI risks inside existing Part 500 controls, including access privileges and multi-factor authentication, rather than adding new requirements.[22]

Know where the sensitive data goes.

AI multiplies copies of data: prompts, retrieval indexes, caches, evaluation sets, logs, fine-tuning sets and provider-side retention. Each copy needs a classification, an owner and a deletion rule.

DATA PATH	CONTROL	LENS
Into a provider's model	Approved route per data class; contract terms on retention and training use; region	Legal, cyber
Into a retrieval index	Labels carried from source; access filtered before ranking; deletion propagated	Cyber, risk
Into logs and traces	Minimum necessary content; masking; retention class; access by role	Legal, cyber
Into evaluation and fine-tuning sets	Training export gate; consent and purpose check; de-identification where possible	Legal, risk
Into memory	Declared readers, writers and retention; erase on request	Legal, cyber

Where data security posture management fits

Data security posture management (DSPM) tools scan stores for sensitive data and risky exposure. They help discovery and they catch copies nobody declared, which is why section 04 lists them as a source of shadow AI. They are not the enforcement point. Treat a DSPM finding as a proposal to the data owner to classify or remove, and keep enforcement in the data layer and the gateway, where the decision is made on each request.

Privacy obligations this serves

These controls produce the evidence that privacy law asks for. Under GDPR, controllers carry out a data protection impact assessment for high-risk processing (Article 35), use processors only on terms in Article 28, and notify the supervisory authority of a personal data breach within 72 hours where feasible (Article 33).[41] Under California's regulations in force since 1 January 2026, businesses that process personal information in ways that present significant risk must conduct risk assessments, and duties for automated decision-making technology used for significant decisions apply from 1 January 2027.[16]

You can outsource the model, not the accountability.

Almost every AI product depends on someone else's model, platform, tools or data. US bank guidance on third-party relationships applies to AI providers like any other: plan, diligence, contract, monitor and plan the exit.[3] The OCC proposed principles-based replacement guidance in September 2026; the 2023 guidance applies until a final version is issued.[4]

SUPPLIER TYPE	PARTICULAR RISKS	WHAT TO CHECK BEFORE RELIANCE
Hosted model providers	Silent model updates; data retention and training use; regional processing; quota and credit limits; outages	Version pinning or change notice; data terms per route; region; rate and spend limits; status history
Cloud AI platforms	Shared responsibility gaps; default logging of prompts; concentration on one cloud	Configuration baseline; logging and retention settings; exit and portability
Open-weight models	Provenance; license terms; unknown training data; you own patching and hosting	Source and checksum; license review; your own evaluation and red team
MCP servers and plugins	Broad scopes; changed behavior after install; token passthrough; untrusted publishers	Owner, pinned digest, scopes, side effects; security test pack; refuse unregistered or changed tools
AI features in licensed software	Switched on by the vendor; data shared with sub-processors; no visibility of the model	Feature inventory; contract addendum; admin controls; registry record
Data and enrichment vendors	External consumer data used in insurance or credit decisions; unclear rights	Data rights; fairness review; NYDFS Circular Letter 7 where it applies[18]

The registry's component list doubles as an AI bill of materials: for each product, every model, prompt, tool, server and data product with its version. When a supplier announces a vulnerability or a change, the registry answers which products are affected in minutes.

Concentration and fallback

If one provider serves many products, its outage, price change or policy change is a firm-wide event. EU financial entities must manage ICT third-party risk as part of their ICT risk framework and remain fully responsible for their obligations when they use providers (DORA Article 28).[39] Record which products share each provider in the registry, so the concentration is visible and the fallback plan is per product, not per hope.

A fallback that was never tested is a second incident.

Consider a generic case. Northwind Mutual, a fictional insurer, runs a claims assistant on one model provider. One morning the provider account reaches its prepaid credit limit and every call is refused. Claims handlers see errors; queues grow.

The team switches the assistant to a second provider it already uses elsewhere. Answers come back, and the incident looks closed. Two weeks later a quality review finds the second model summarized policy exclusions less completely, and its answers cited sources in a format the output check did not parse, so ungrounded statements passed. Nobody had evaluated that route for this product. The fix for an outage became a quality incident.

LESSON	CONTROL
Credit, quota and spend are availability risks	Monitor provider credit, quota and rate limits as service signals with alerts and owners, not only as finance data
A fallback is a release	Register every fallback route as a component of the product; evaluate it against the same criteria and critical-failure veto before it can serve
Data terms differ by provider	Approve each route for the data classes it may receive; an unapproved route gets no restricted data and no tools
Switching is a change	A provider switch under incident is an emergency change: approved by a named person, logged, and reviewed afterwards
Degraded mode can be better than a substitute	Where no evaluated fallback exists, return a defined incomplete state or route work to people

Contract terms that support the control

Contract terms that support these controls include notice of model changes, incident notification within stated hours, sub-processor lists, data location, evidence and audit access, assistance during incidents, and exit with data return. For EU financial entities, DORA Article 30 lists required contractual provisions, including service locations, data access and return on termination, incident assistance and cooperation with authorities.[39] Section 12 covers the liability and intellectual property terms.

OWNER The first line owns the supplier relationship and the fallback test. Second-line risk owns the concentration view. The CISO office reviews security terms and the supply chain checks.

US federal: banking supervision, the FTC and federal policy.

This part maps obligations to registry fields and controls, so legal duties become things the framework records and tests. It is a map for counsel to confirm, not a statement of every duty. Supervisory detail is in AI Governance Evidence.

OBLIGATION AND STATUS	REGISTRY FIELDS	CONTROLS AND EVIDENCE
SR 26-2, model risk guidance, 17 April 2026: supersedes SR 11-7; nonbinding; generative and agentic AI outside scope, left to banks' broader governance[1]	Model or not; governing policy; validator	Inventory, tiering, independent validation in the gate
OCC Bulletin 2026-54, examiner framework, 4 November 2026: six pillars; sets no new expectations for banks[2]	Purpose statement; tier; pillar tags on evidence	Evidence pack indexed by purpose, risk, control, oversight, resilience, assurance
Third-party relationships guidance, June 2023; proposed replacement 11 September 2026[3, 4]	Vendor record; criticality; exit plan	Due diligence, contract review, monitoring, tested fallback
Regulation B adverse action notices: specific principal reasons[6]	Decision type; reason codes	Reasons traceable to the request trace; tested explanations
Computer-security incident notification: notify the primary federal regulator within 36 hours of determining a notification incident[5]	Incident record; clock	Incident playbook with the clock table (section 13)
FTC Act Section 5: deceptive AI claims; Operation AI Comply, 25 September 2024[9]	Marketing and disclosure claims	Claims about AI capability reviewed against evaluation evidence

FTC posture has shifted. In December 2025 the FTC set aside its 2024 Rytr order, and in July 2026 it proposed a policy statement on suppression of accuracy in AI systems, which is not final.[10, 11] The constant is the deception standard: do not claim more for an AI product than your evidence shows.

Executive Order 14365 of 11 December 2025 seeks a national framework and directs challenges to state AI laws. An executive order does not itself repeal a state law, so keep state duties in the map until counsel says otherwise.[8]

Sector law: financial privacy, health and California.

Sector privacy and security laws apply to AI exactly as they apply to any system that touches the data. The control map makes that explicit for each product.

OBLIGATION AND STATUS	REGISTRY FIELDS	CONTROLS AND EVIDENCE
Gramm-Leach-Bliley Act: privacy of nonpublic personal information and safeguards[12]	Data classes reached; route approvals	Masking, purpose binding, approved routes per data class, provider contract terms
FTC Safeguards Rule, for financial institutions under FTC jurisdiction: notify the FTC within 30 days of discovering an event involving 500 or more consumers, effective 13 May 2024[13]	Applicability; incident clock	Incident playbook; provider incident notice in contract
HIPAA privacy, security and breach rules: individual breach notice within 60 days of discovery[14]	PHI flag; business associate terms	Approved routes for PHI; agreements with providers that handle PHI; breach clock
California regulations on risk assessments and automated decision-making technology, effective 1 January 2026; ADMT duties for significant decisions from 1 January 2027[16]	California consumers; ADMT use; significant decision	Risk assessment record; notice and access process; evaluation evidence

Healthcare note

Health insurers, and banks and insurers that process health data, add two layers. HIPAA governs protected health information wherever an AI product reads, stores or sends it, including to a model provider acting as a business associate.[14] HHS's 2024 rule under Section 1557 of the Affordable Care Act requires covered entities to make reasonable efforts to identify patient care decision support tools that use race, color, national origin, sex, age or disability as inputs, and to mitigate the risk of discrimination, with compliance from 1 May 2025 as published; counsel should confirm current status.[15]

For the framework this means three registry flags: whether a product touches protected health information, whether it supports clinical or coverage decisions, and which inputs it uses that are protected characteristics. Each flag raises the tier and adds evidence: an approved route for health data, a fairness review and a named clinical or medical owner. Workflow Studio's payer prior authorization template is an example of a workflow where a person decides and the AI prepares.

STATE PRIVACY LAWS — Several other states have comprehensive privacy laws with profiling and risk assessment duties. Record each state where the product affects residents in the jurisdictions field, and let counsel maintain the list.

US states: insurance regulators, Colorado and New York.

Insurance is regulated state by state, and New York regulates cybersecurity across banking and insurance. These are the obligations most likely to apply to a US insurer or a New York-licensed firm.

OBLIGATION AND STATUS	REGISTRY FIELDS	CONTROLS AND EVIDENCE
NAIC Model Bulletin on insurers' use of AI, December 2023: a written AI systems program, governance, risk controls, third-party oversight, testing for adverse consumer outcomes; adopted state by state, see the NAIC map[17]	Adopting states where the product is used; consumer decision	Registry as program inventory; gate evidence; vendor records
NYDFS Insurance Circular Letter No. 7, July 2024: governance and fairness testing for AI and external consumer data in underwriting and pricing[18]	External data sources; underwriting or pricing use	Data source approval; fairness tests in the gate
Colorado Regulation 10-1-1: governance framework, inventory, board oversight and reports for life, private auto and health benefit plan insurers[19]	Line of business; external data; inventory	Inventory export; committee reporting; vendor oversight
Colorado SB 26-189, signed 14 May 2026: repeals and replaces SB 24-205; scheduled to apply from 1 January 2027; confirm enforcement status with counsel[20]	Colorado consumers; automated decision	Track status; prepare notice and evidence
23 NYCRR Part 500, as amended 1 November 2023: risk assessment, access, monitoring; notice within 72 hours of determining a cybersecurity incident, 24 hours for an extortion payment[21]	NY-licensed entity; incident clock	Threat model; access tests; incident playbook
NYDFS industry letters on AI, 16 October 2024 and 21 May 2026: AI risks inside existing Part 500 controls; frontier AI threat measures; no new requirements[22, 23]	Threat model; third parties	Risk assessment update; third-party coordination; review of AI-generated code

Check each state where you write business. Several large states, including New York and Colorado, have their own rules instead of the NAIC bulletin, and adoption keeps changing.

PROGRAM, NOT PROJECT Where the NAIC bulletin applies, regulators expect an insurer-wide AI program. The registry, the gate and the committee report are that program's evidence.

European Union: the AI Act, DORA and GDPR.

Firms with EU operations or EU customers add three regimes. The same registry fields carry them.

OBLIGATION AND STATUS	REGISTRY FIELDS	CONTROLS AND EVIDENCE
AI Act classification: Annex III lists creditworthiness assessment of natural persons and risk assessment and pricing in life and health insurance; Article 6(3) narrow exceptions must be documented[37]	EU role (provider, deployer); Annex III use; Article 6(3) rationale	Applicability record approved by counsel
Timetable after Regulation (EU) 2026/1744 (in force 27 July 2026): Annex III high-risk rules from 2 December 2027; high-risk systems in regulated products from 2 August 2028; Article 50(2) marking by 2 December 2026 for generative systems already on the market[38]	Applicable dates per product	Dated applicability record; review trigger
AI literacy (Article 4, as amended): measures to support staff AI literacy[38]	Training owner	Role-based training records
Deployer duties for high-risk systems (Article 26): competent human oversight; monitoring; logs kept at least six months unless other law provides; financial institutions meet some duties through existing governance. Deployers of credit scoring and life and health insurance pricing systems also carry out a fundamental rights impact assessment (Article 27)[37]	Oversight owner; log retention class; impact assessment	Oversight exercise; retention schedule; assessment record
Serious incidents (Article 73): providers report within 15 days of awareness; 10 days on a death; 2 days for widespread infringement or critical infrastructure[37]	Incident clock	Playbook; incident record
DORA: ICT third-party risk (Article 28), contract terms (Article 30), major ICT incident reports: initial within 4 hours of classification and 24 hours of awareness, intermediate within 72 hours, final within one month[39, 40]	Register of information; criticality	Vendor records; incident playbook
GDPR: solely automated decisions with legal or similar effect (Article 22); DPIA (Article 35); breach notice within 72 hours where feasible (Article 33)[41]	Personal data; automated decision	DPIA in the gate; oversight design; breach clock

EARLIER DATES STILL APPLY A later high-risk date does not excuse duties that already apply, such as AI literacy and the prohibited practices. AI Governance Evidence has the full timetable.

United Kingdom, Canada, Singapore and global standards.

Outside the US and the EU, the main supervisors of banks and insurers have chosen to apply existing model risk, conduct and accountability rules to AI rather than write a separate AI statute. The registry serves them with the same records.

OBLIGATION AND STATUS	REGISTRY FIELDS	CONTROLS AND EVIDENCE
UK PRA SS1/23, model risk management principles for banks, effective 17 May 2024: inventory, tiering, independent validation and governance, including AI-based models[42]	Model inventory entry; model tier; validator	Independent validation in the gate; registry as inventory
UK FCA approach: no AI-specific rules; relies on existing frameworks including the Consumer Duty and the Senior Managers and Certification Regime, under which senior managers are accountable for their firm's use of technology and AI[43]	Accountable senior manager; consumer outcome	Named senior manager per product; consumer outcome monitoring
Canada, OSFI Guideline E-23, final 11 September 2025, effective 1 May 2027: model risk management for all federally regulated financial institutions, explicitly including AI and machine learning models[44]	Model inventory; risk rating; third-party model	Inventory, risk rating and lifecycle controls; third-party models under the same standard
Singapore, MAS Guidelines on AI Risk Management, issued 7 October 2026: board and senior management oversight, AI inventories, risk materiality assessment, proportionate lifecycle controls and accountability for AI provided by third parties[45]	Materiality; third-party AI	Registry as inventory; tier as materiality; vendor controls
Financial Stability Board sound practices, final 15 October 2026: twelve voluntary practices, including triggers for agent human approval and tested ways to stop[7]	Autonomy per action; kill switch	Approval binding; kill-switch test
ISO/IEC 42001:2023, AI management system: certifiable, voluntary[35]	Management system scope	Registry and gate as the operational core of the system

The pattern across regimes is consistent: know what AI you use, rate it, control it in proportion, keep a named person accountable, and be able to show the evidence. One registry built well answers all of them; a program per regulator multiplies the work and the inconsistencies.

When an agent acts, the firm acted.

Treat an action taken by an AI agent under your firm's credentials as an action of the firm. Regulators already take that view of outsourcing: EU financial entities remain fully responsible for their obligations when they use ICT providers, Singapore's guidelines keep institutions accountable for AI provided by third parties, and UK senior managers are accountable for their firm's use of AI where it affects regulated activities.[39, 45, 43] Counsel should advise on agency and contract law in each jurisdiction; the governance answer is the same everywhere: bound the authority, bind the approval, keep the record.

CONTRACT TERM	WHY IT MATTERS	WHERE IT LANDS IN THE REGISTRY
Use of your data for training	Prompts and outputs may contain confidential and personal data	Vendor record: training use prohibited or opt-in
Data location and sub-processors	Privacy and residency duties; DORA location terms	Vendor record; route approval per data class
Model change notice	A silent update is an unreviewed change	Change trigger; version pin
Security and incident notice in hours	Your own clocks start when you know	Incident clock table
Evidence and audit access	You must show controls work, including the vendor's part	Evidence pack; audit plan
Warranties on performance	Few providers warrant output accuracy; know what is and is not promised	Gate evidence must come from your own evaluation
Intellectual property indemnity	Claims that outputs or training infringe third-party rights	Legal record; residual risk accepted by counsel
Liability caps and exclusions	Caps often sit far below the harm an AI error can cause	Residual risk entry; insurance review
Exit, data return and deletion	Concentration and continuity	Exit plan; tested fallback

Privilege, confidentiality and insurance

Keep privileged and highly confidential material on routes whose terms give the provider no access for training or review, and record that on the route. Check whether existing cyber, errors and omissions and directors' policies cover losses caused by AI errors and agent actions; record the answer as part of the residual risk accepted at the gate.

PRACTICAL RULE Do not rely on a vendor warranty as release evidence. Use contracts to allocate risk and secure cooperation; use your own evaluation to decide whether to rely on the product.

Output rights, records and disclosure.

Three legal questions recur for every AI product: who owns what it produces, what must be kept, and what must be told.

Training data and output rights

The US Copyright Office concluded in January 2025 that copyright protects AI-assisted work only where a human author determined sufficient expressive elements, and that prompts alone do not make the user an author.[36] Do not assume your firm owns every AI output in the way it owns staff work product. Allocate rights in contracts, keep human authorship where ownership matters, and record in the registry whether a product generates content that is published or sold. For models you fine-tune or train, record the source and license of every training set.

Records and legal hold

Prompts, retrieved passages, outputs, approvals and agent actions can be business records and discoverable evidence. Give each record type a retention class approved by counsel, keep the trace that explains a decision for as long as the decision can be challenged, and make legal hold suspend deletion across every copy: logs, indexes, caches and provider-side storage where the contract allows. For high-risk systems under the EU AI Act, deployers keep logs for at least six months unless other law provides otherwise, and financial institutions keep them within their existing documentation.[37]

Disclosure

DUTY	CONTROL
Tell people when they deal with AI or receive AI-generated content where law requires it (for example AI Act Article 50)[37]	Approved notice wording and placement, tested in the interface
Give specific reasons for credit decisions (Regulation B)[6]	Reasons traceable to the trace; tested explanations
Do not overstate what the AI does (FTC Act Section 5)[9]	Marketing claims reviewed against evaluation evidence
Notices and rights for automated decision-making (California from 2027; GDPR Article 22)[16, 41]	Applicability record; notice and request process

OWNER Legal and compliance own the retention schedule, the notice wording and the rights allocation. The first line implements them; audit tests that legal hold actually stops deletion.

Watch what was approved; know an incident when you see one.

Monitoring answers one question continuously: is the product still behaving within the conditions of its approval? The thresholds come from the gate. Breaching one is a decision point, not just a chart.

SIGNAL	EXAMPLE THRESHOLD	LENS	FIRST OWNER
Quality: grounding, critical failures, incomplete-answer rate	Any critical failure; groundedness below the approved floor	Risk	Product owner
Human oversight: review rate, override rate, reviewer backlog	Overrides above the approved band; backlog past due time	Risk	Business owner
Fairness: outcome differences across groups where lawful to measure	Difference beyond the approved tolerance	Legal, risk	Business owner
Security: blocked injections, tool denials, unusual data volume, unknown release IDs	Spike against baseline; any unknown release	Cyber	CISO office
Cost and capacity: spend, quota, credit, rate-limit errors	80 percent of budget or credit; sustained errors	Risk	Product owner
Drift and change: provider version, source freshness, prompt or tool digest	Any unapproved component change	All	Product owner
Complaints and feedback	Any complaint citing the AI	Legal	Business owner

Defining an AI incident

An AI incident is an event in which an AI product caused, or could reasonably have caused, harm to customers, staff, the firm or third parties, or acted outside its approval. That includes a disclosure, a wrong consequential decision, an action outside the agent contract, a successful injection, a breach of a fairness tolerance, an unapproved component in production, or loss of service on a critical workflow. Use the firm's existing incident process and severity scale; add the AI-specific triggers and the registry ID to it. A separate AI incident process creates a second clock that someone will miss.

SEVERITY Severity follows the tier and the harm, not the technology. A Tier 1 product acting outside its contract is at least high severity even if no loss is yet known.

Clocks, kill switches and postmortems.

One AI incident can start several clocks at once. Map them before the incident, per product, in the registry. Each clock runs from the moment the regime says, which is often determination or awareness, not the end of the investigation.

REGIME	WHO AND WHEN	CLOCK
US banking agencies[5]	Bank to primary federal regulator, after determining a notification incident	36 hours
NYDFS Part 500.17[21]	Covered entity to the superintendent, after determining a cybersecurity incident; extortion payment	72 hours; 24 hours
FTC Safeguards Rule[13]	Non-bank financial institution to the FTC, 500 or more consumers, after discovery	30 days
HIPAA[14]	Covered entity to individuals after discovering a breach	Within 60 days
GDPR Article 33[41]	Controller to supervisory authority, where feasible	72 hours
DORA[40]	Initial report after classifying a major ICT incident, and at the latest from awareness	4 hours; 24 hours
AI Act Article 73[37]	Provider of a high-risk system, after awareness of a serious incident (high-risk rules apply from 2027)	15, 10 or 2 days
Bank service providers[5]	Provider to its bank customers, for an incident that disrupts covered services for four or more hours	As soon as possible
Contracts and state laws	Customer, partner and state breach notification duties	Per contract and state

Stop first, then learn

- **Kill switches** per agent, tool, product and route, operated by the first line, effective within a stated time, tested at the gate and twice a year.
- **Containment** keeps the evidence: freeze the release, keep the trace, do not alter the system in ways that destroy evidence before the investigation.
- **Postmortem into the registry**: root cause, the control that failed or was missing, the change approved to fix it, and whether the tier or the evaluation set changes. The fix goes back through the gate.

ACTIVITY	BUSINESS OWNER (1L)	PRODUCT AND TECH OWNER (1L)	RISK AND MODEL RISK (2L)	LEGAL AND COMPLIANCE (2L)	CISO OFFICE (2L)	INTERNAL AUDIT (3L)
Operate monitoring and respond to alerts	**A**	R	I	I	C	T
Declare an AI incident and set severity	**A**	R	C	C	R	I
Pull the kill switch	**A**	R	I	I	C	I
Decide and make regulatory notifications	C	C	C	**A**	R	I
Postmortem and registry update	**A**	R	C	C	C	T

R does the work · A accountable for the outcome · Ap approves independently · C consulted · I informed · T tests afterwards. One A per row. 1L, 2L, 3L: first, second and third line.

Challenge in the second line, testing in the third.

Oversight works when each line can see the same facts and has a defined job to do with them. The registry supplies the facts; this section defines the jobs.

Second-line challenge

- Review every Tier 1 and Tier 2 release before approval, with challenge recorded on the record.
- Sample Tier 3 and Tier 4 releases monthly and re-tier where needed.
- Review monitoring breaches, open exceptions and incidents against risk appetite, and escalate.
- Own the standards: tiering method, evaluation minimums, red-team scope, control map, retention classes.

Third-line testing

Internal audit plans AI work under the IIA's Global Internal Audit Standards, effective January 2025, using the registry as the population.[34] Typical tests:

AUDIT TEST	METHOD	WHAT A FINDING LOOKS LIKE
Completeness	Compare discovery sources with the registry	AI in use without a record
Gate operation	Re-perform the gate on a sample: tier, evidence, approvers, independence	Approval without required evidence; builder as approver
Runtime enforcement	Call the gateway with an unapproved release, a revoked user, an out-of-scope tool	A call that should be denied is allowed
Exceptions	Trace exceptions from request to closure or expiry	Expired exception with the product still running
Incident handling	Follow incidents through clocks, notifications and postmortems	Missed clock; fix not re-gated
Change	Compare deployed components with approved release IDs	Unapproved model or prompt version in production

EXAMINER-READY EVIDENCE

Tag each artifact in the registry by the OCC framework's six pillars: purpose, risk, control, oversight, resilience and assurance.[2] When an examiner asks a pillar question, the answer is a filtered list from the registry, not a new document. AI Governance Evidence lists what reviewers accept as evidence for each.

One report to the board, built from the registry.

Board committees need few numbers, consistently defined, with trends. Build the report from the registry so that it cannot drift from what is deployed.

METRIC	DEFINITION	WHY THE COMMITTEE CARES
Inventory by tier	Products in production by tier, and change since last quarter	Scale and direction of exposure
Unregistered AI found and closed	Discovery findings, and those registered or switched off	Whether the registry is the real door
Gate outcomes	Releases approved, declined and sent back, by tier; median time to decision	Whether the gate says no, and whether it is a bottleneck
Exceptions	Open exceptions by tier and age; expired and still running	Risk carried outside appetite
Incidents	AI incidents by severity; notifications made; repeat causes	Harm and learning
Control health	Share of Tier 1 and 2 products with all runtime controls enforcing, not warning	Approval matched by enforcement
Human oversight	Override and review rates against approved bands	Whether people still exercise judgment
Material changes	Changes re-gated; unapproved components detected	Change control working
Audit findings	Open findings by rating and age	Independent view

The risk committee owns appetite and the second-line view; the audit committee owns the third-line view. Both see the same registry figures. Colorado's insurance regulation and Singapore's new guidelines both expect board-level oversight of AI risk; this report is how it becomes routine.[19, 45]

Between reports, the second line escalates any Tier 1 incident, any expired Tier 1 exception and any declined Tier 1 release to the committee chair within days, not at the next meeting.

FROM PRACTICE

In my experience, committees engaged once the report showed declined releases and expired exceptions, not only counts of use cases. A gate that never declines anything tells a board nothing. A gate that declined three releases last quarter, and shows why, tells it the control is real.

Decide which changes go back through the gate.

Most AI risk enters after launch. A provider updates its model, someone edits a prompt, the index gains a new document set, an agent receives a tool. Define material change in advance so the decision is not made under deadline pressure.

CHANGE	MATERIAL WHEN	WHAT RE-TESTING IS NEEDED
Model version or provider	Always for Tier 1 and 2; for Tier 3 and 4 when the route or provider changes	Full workflow evaluation, critical-failure veto, red-team gateway cases, data terms check
Fallback route activated	Always, unless the route already passed the same evaluation	As for a model change, before use; under incident, an emergency change reviewed afterwards
Prompt or system instruction	Any change to a Tier 1 or 2 prompt; behavior-affecting changes otherwise	Evaluation pack run; approval by someone other than the author
Tool added or scope widened	Always	Threat model update; tool tests; autonomy review; re-tier
Retrieval corpus or data product	New source, new classification, or new purpose	Access tests; grounding evaluation; privacy review
Users, jurisdiction or purpose	New user group, country or decision type	Re-tier; legal control map update
Autonomy level	Any increase	Full gate at the new level

A change that is not material still needs a record: who changed what, when, and why it was judged not material. The second line samples those judgments. EU law makes some changes legally significant as well; under the AI Act a substantial modification can change who carries provider duties, so counsel should review changes to high-risk systems.[37]

Emergency changes happen. Allow them under a named approver, log them at the time, and bring them back through the gate within a set number of days. An emergency change that never returns to the gate is an exception without an end date.

MODEL OR PROVIDER SWAP	Treat a swap as a new release of every product that uses the model. The registry lists them; the evaluation packs already exist; the work is to run them, read them and approve or decline.

Retire with records, not just a shutdown.

Retirement is a control point because AI leaves things behind: credentials, indexes, cached data, fine-tuned weights, agent memory, open approvals and records that must be kept.

RETIREMENT STEP	EVIDENCE
Decide and communicate	Retirement decision on the registry record, with users and dependent products notified
Revoke identities and credentials	Agent identity disabled; tokens and keys revoked; tool scopes removed
Remove data copies	Indexes, caches, memory and evaluation sets deleted or archived under their retention class; provider deletion requested and confirmed where the contract allows
Keep records	Trace, approvals, incidents and evaluation results kept for their retention period; legal hold respected
Close commitments	Contracts ended or reduced; exceptions closed; monitoring switched off last
Confirm	Gateway refuses the retired release ID; audit can sample the closure

ACTIVITY	BUSINESS OWNER (1L)	PRODUCT AND TECH OWNER (1L)	RISK AND MODEL RISK (2L)	LEGAL AND COMPLIANCE (2L)	CISO OFFICE (2L)	INTERNAL AUDIT (3L)
Propose a change and judge materiality	**A**	R	C	C	C	I
Approve a material change, Tier 1 and 2	**A**	R	**Ap**	C	**Ap**	T
Run re-evaluation after a model or provider swap	I	**A**	C	I	C	T
Sample non-material change judgments	I	I	**A**	C	C	T
Decide retirement	**A**	R	I	C	I	I
Revoke access and remove data copies	I	**A**	I	C	R	T
Approve record retention and legal hold	I	R	I	**A**	I	T

R does the work · A accountable for the outcome · Ap approves independently · C consulted · I informed · T tests afterwards. One A per row. 1L, 2L, 3L: first, second and third line.

Periodic review closes the loop. Each product has a next review date by tier: at least annually for Tier 1 and 2. A missed review date is treated like an expired exception: the product is flagged, and the committee sees it.

How RJL's three products carry the controls.

RJL runs the patterns in this paper in three products on AWS: Ontology Studio for approved business meaning, Workflow Studio for reviewed workflows where a person decides, and the Context Engine for governed context, the registry and the evidence trail. Statuses below come from RJL's product records of 20 November 2026: **Live** means deployed and checked live. **Built, beta** means built and tested, in beta while we complete live checks. **Available, off by default** means built and ready; switched on per client or demo. The examples run on fictional data.

CONTROL	PRODUCT	FEATURE	STATUS
Registry and inventory	Context Engine	AI inventory with discovery check and sealed export	LIVE
Registry and inventory	Context Engine	Agent lifecycle with gated promotion, stop, degrade and retire	LIVE
Registry and inventory	Context Engine	Prompt registry: approved versions only, promotion needs an evaluation and an independent approver	BUILT, BETA
Tiering	Context Engine	Tier per use case on one named scale	BUILT, BETA
Gate and independence	Context Engine	Decision-rights register with independence checks applied by every workflow	LIVE
Gate and independence	Context Engine	Release decision records signed by three different people	BUILT, BETA
Gate and independence	Context Engine	Release ID and policy hash on every audit record; changed composition without a passed evaluation serves no agent	LIVE
Gate and independence	Workflow Studio	Saved test cases and a release gate that refuses until tests pass	LIVE
Three lines	Context Engine	Every role carries its line; internal audit read-only; second-line approval for critical items	BUILT, BETA
Exceptions	Context Engine	Exceptions with owner, compensating control, expiry and retest	BUILT, BETA
Legal map	Context Engine	Applicability records confirmed by the second line	BUILT, BETA
Runtime guardrails	Context Engine	Governed gateway: policy decision before any read on every call	LIVE
Runtime guardrails	Context Engine	Tool registry with pinned digests; unregistered or changed tools refused	LIVE
Runtime guardrails	Context Engine	Every tool result screened for planted instructions	LIVE
Runtime guardrails	Context Engine	Run limits and spend caps per agent with defined fallbacks	LIVE
Identity	Context Engine	Delegated token per answer naming person, agent, audience and purposes	BUILT, BETA

What runs, what is in beta, and what is not there.

CONTROL	PRODUCT	FEATURE	STATUS
Identity	Context Engine	Approval bound to the exact change with expiry and single use	BUILT, BETA
Data protection	Context Engine	Tokenization of restricted identifiers, two-person detokenize	AVAILABLE, OFF BY DEFAULT
Data protection	Workflow Studio	Personal and patient data check before any AI call	LIVE
Evaluation	Context Engine	Negative release suite on every packaging run	LIVE
Evaluation	Context Engine	Red-team library generated from policy	BUILT, BETA
Third parties	Context Engine	Approved-model register with data terms per route	BUILT, BETA
Third parties	Context Engine	Supplier records with exit plans and a portability drill	BUILT, BETA
Runtime and incidents	Context Engine	Kill switches per agent, tool, product and role	LIVE
Runtime and incidents	Context Engine	Incident records, reporting clocks per regime, postmortems	LIVE
Evidence	Context Engine	Keyed, hash-chained audit of every call with nightly verification	LIVE
Evidence	Context Engine	Retention classes per record and legal hold	LIVE
Assurance	Context Engine	Examiner portal: read-only examiner access and evidence pack by OCC pillar	BUILT, BETA
Meaning	Ontology Studio	Drafts arrive as labelled proposals that never self-approve	LIVE
Human review	Workflow Studio	Reviewer inbox: accept, return or reject	LIVE

What this shows, honestly

Several second-line controls are built but not yet enforcing: the three-lines roles run in warn mode with no auditor accounts yet, the release decision gate runs in warn mode, and the exceptions workbench waits for its sign-in groups. That is the state many firms will recognize: the design exists, the enforcement switch is the decision. RJL's own model routing is an example of section 10's rule. RJL's own fallback route was evaluated on 20 November 2026 and did not meet the threshold; fixes are being made and the route will be re-tested before it is relied on.

Not supported today: enforcement inside a customer's own data platform, an AI operations triage agent, and a confidential computing runtime. The full feature list, with links where a feature can be seen working, is at rjlsol.com/reference-implementation.

Thirteen weeks to a working registry and gate.

Start with the door, not the policy. In ninety days a bank or insurer can have one registry that deployment reads, one questionnaire, a working gate for one Tier 2 product and its first committee report.

WEEKS	DELIVERABLE	LINES INVOLVED	EVIDENCE AT THE END
1 to 2	Name the owners: executive sponsor, registry owner, second-line approvers by lens, audit contact. Agree the tiering standard in draft	All three	Signed charter; draft standard
3 to 4	Stand up the registry with the minimum record (Appendix A). Load known AI from the three existing review lists	First, second	Registry with every known product and owner
5 to 6	Run discovery (section 04). Register or schedule switch-off for each finding. Adopt the questionnaire (Appendix B)	First, second	Unregistered AI count; first tiers
7 to 8	Pick one Tier 2 product. Assemble its evidence by tier. Write its threat model and legal control map entry	First, second	Complete evidence pack
9 to 10	Run the gate for real: independent approvals, conditions, expiry. Connect deployment to the registry release ID	All three	Signed release; pipeline refuses unapproved IDs
11	Turn on runtime checks for that product: gateway reads the release ID; monitoring thresholds from the approval; kill-switch test	First, CISO office	Denied-call samples; kill-switch result
12	Map incident clocks into the playbook (Appendix E); run a tabletop on one AI incident	First, second	Tabletop record and fixes
13	First committee report from the registry. Internal audit walkthrough of the gate on the pilot	All three	Committee minutes; audit memo

After ninety days, extend the gate to every Tier 1 and Tier 2 product on a published schedule, with existing products brought through in order of risk. The 90-Day First Governed Capability describes the delivery gates in more detail, and the Production-Readiness Diagnostic gives a ten-question check of where you stand.

WHAT SUCCESS LOOKS LIKE By day ninety, a release without a registry record cannot reach production, one product has passed a real gate, and the board has seen the first report. Everything else is scale.

APPENDIX A / REGISTRY RECORD TEMPLATE

The registry record, field by field.

Required before the gate unless marked. Field names are suggestions; keep them stable once chosen, because runtime, reporting and audit all read them.

FIELD	CONTENT	SET BY	REQUIRED FOR
registry_id	Stable identifier for the product, never reused	Registry	All
release_id	Identifier for one approved composition; changes when any component changes	Registry	All
name, description	Plain-language name and one-paragraph description	First line	All
business_owner	Named person accountable for the use and its outcome	First line	All
product_owner	Named person accountable for build and operation	First line	All
senior_manager	Accountable executive where a regime names one (for example SM&CR)	First line	Tier 1 and 2
purpose	Intended use, decisions informed, out-of-scope uses	First line	All
users, affected_people	Who uses it; who is affected, how many, whether consumers	First line	All
autonomy	Level L0 to L4 for each action the product can take	First line	All
components	Models and versions, prompts, tools and MCP servers, data products, vendors, each with its own registry entry	First line	All
routes	Model routes with approved data classes, including fallback routes	First line, CISO office	All
data_classes	Classes read, written and sent to providers; PII, NPI, PHI flags	First line	All
jurisdictions	Countries and states where it is used or affects residents	First line, legal	All
questionnaire	Answers to Appendix B, versioned	First line	All
tier, tier_rationale	Proposed tier and why; confirmation for Tier 1 and 2	First line; second line confirms	All
legal_map	Applicable obligations from section 11, with registry fields and controls	Legal and compliance	Tier 1 to 3
threat_model	Threats from section 07 that apply, with controls	First line; CISO office signs	Tier 1 to 3

Release, operation and closure fields.

FIELD	CONTENT	SET BY	REQUIRED FOR
evidence	Links to each evidence item required for the tier (section 06), with version and date	First line	All
evaluation	Pack version, criteria set before testing, results, critical failures, evaluator	First line; independent for Tier 1	All
red_team	Plan (Appendix D), findings, fixes, residual risk	First line; CISO office	Tier 1 to 3
impact_assessments	Privacy impact assessment, DPIA, fundamental rights or other assessments required	Legal and compliance	As triggered
vendor_review	Contract terms, sub-processors, incident notice, exit plan, concentration flag	First line; risk	Where a vendor is used
approvals	Named approvers, line, date, scope, conditions, expiry; recorded challenge	Approvers	All
exceptions	Each with owner, compensating control, acceptor, start and end dates	First line; second line accepts	As needed
conditions	User population, allowed actions, data boundary, thresholds that trigger review	Approvers	All
monitoring	Signals, thresholds, owners, dashboards	First line	All
kill_switch	Mechanism, operator, last test date and result	First line	Tier 1 to 3
fallback	Degraded mode and evaluated fallback routes, with test date	First line	Tier 1 to 3
incident_clocks	Regimes and contracts that apply, with clock and recipient	Legal and compliance	Tier 1 to 3
retention_class	Retention per record type; legal hold flag	Legal and compliance	All
changes	Change log with materiality judgment and re-gate reference	First line	All
incidents	Incident IDs, severity, postmortem link	First line	All
occ_pillar_tags	Pillar tag on each evidence item, for US banks	First line; second line checks	US banks
next_review	Date of next periodic review by tier	Second line	All
status	Proposed, approved, live, restricted, retired	Registry	All
retirement	Decision, revocations, data removal, records kept, confirmation	First line	At retirement

KEEP IT SMALL Start with the fields marked All. The registry earns trust by being complete and current, not by being long.

APPENDIX B / RISK ASSESSMENT QUESTIONNAIRE

One questionnaire for three lenses.

Answer every question. Score each 0 to 3 as shown. Hard triggers set a minimum tier whatever the score.

#	QUESTION	SCORING (0 / 1 / 2 / 3)	HARD TRIGGER
U1	What does the output inform or do?	Internal information / internal work product / decision support for a business decision / decision or action affecting a customer, money or a record	U1 = 3 sets at least Tier 2
U2	How reversible is an error, and how fast is it detected?	Trivially / easily, within a day / with effort / hard or impossible	
A1	Highest autonomy level of any action	L0 / L1 / L2 / L3 or L4	L3+ with U1 = 3 sets Tier 1
A2	Can it act with no person reviewing that action?	No / on internal data only / on records / on money or customers	
D1	Most sensitive data class it reaches	Public / internal / confidential / NPI, PHI, credentials or regulated	D1 = 3 sets at least Tier 3
D2	Does data leave the firm to a provider?	No / anonymized / pseudonymized / identifiable	
P1	Who is affected?	Staff only, a few / staff at scale / business customers / consumers	
P2	Do affected people know AI is involved, and can they reach a person?	Yes and yes / yes / partly / no	
L1	Does a named legal trigger apply? (credit, insurance pricing or underwriting, claims, employment, AI Act Annex III, automated decisions, health)	None / possible / likely / confirmed by counsel	Confirmed sets at least Tier 2
L2	In how many jurisdictions is it used or does it affect residents?	One / two or three / several / EU plus others	
L3	Is it prohibited or outside AI policy?		Yes: stop
S1	Who can reach it?	Named staff / all staff / partners / public	
S2	Does it read untrusted content (email, web, documents from outside)?	No / curated / mixed / open	
S3	Which tools and outbound channels does it hold?	None / read-only / internal write / external send or payment	S1 ≥ 2 and S3 = 3 sets at least Tier 2

From answers to a tier.

The tiering logic is deliberately simple. It can be applied by hand, checked by a second-line reviewer in minutes, and coded into the registry so the proposed tier is computed, not typed.

STEP	RULE
1 Stop rule	If L3 is yes, the product cannot be registered for production as described
2 Hard triggers	Apply every hard trigger in the table; keep the highest minimum tier
3 Score	Sum the thirteen scored questions (0 to 39)
4 Score to tier	30 to 39: Tier 1. 20 to 29: Tier 2. 10 to 19: Tier 3. 0 to 9: Tier 4
5 Highest wins	The tier is the more severe of step 2 and step 4
6 Judgment	The first line may propose a higher tier for a reason; only the second line may lower a computed tier, with a recorded rationale
7 Confirmation	The second line confirms every Tier 1 and Tier 2 result before evidence work starts

Worked examples

PRODUCT	KEY ANSWERS	RESULT
Claims file summary for adjusters at Northwind Mutual (fictional)	U1 = 2, A1 = 0, D1 = 3, P1 = 3, L1 = 1, S2 = 1; score 16; hard trigger D1 = 3	Tier 3
Agent that drafts and, with approval, sends payment reminders for Contoso Bank (fictional)	U1 = 3, A1 = 2, D1 = 3, P1 = 3, S3 = 3; score 24; hard trigger U1 = 3	Tier 2
Agent that releases fraud holds within limits without per-action review	U1 = 3, A1 = 3; hard trigger L3+ with U1 = 3	Tier 1
Internal policy search for staff	U1 = 0, A1 = 0, D1 = 1, S1 = 1; score 4	Tier 4

Calibrate the thresholds on your own portfolio: tier twenty known products by expert judgment, run the questionnaire, and adjust the score bands until they agree. Record the calibration as part of the tiering standard, and repeat it each year.

WHY ONE QUESTIONNAIRE — Every question is there because one of the three lenses needs it. When risk, legal and security each answer the same record, nobody has to reconcile three versions of what the product does.

APPENDIX C / CONTROL MATRIX

Control, threat, legal source, owner, evidence.

Twenty core controls. Legal sources are the ones in section 11 and the register; counsel confirms which apply to each product. Every control has a first-line operator, a second-line standard owner and third-line testing.

CONTROL	THREAT OR RISK	LEGAL OR SUPERVISORY SOURCE	LINE OF DEFENSE OWNER	EVIDENCE PRODUCED
C01 Registry completeness	Shadow AI; unreviewed use	NAIC bulletin; Colo. 10-1-1; MAS; E-23	1L registers; 2L standard; 3L tests	Registry; discovery reconciliation
C02 Tiering	Under-controlled high-impact use	OCC 2026-54; SR 26-2; MAS materiality	1L proposes; 2L confirms	Questionnaire; confirmation
C03 Release gate and independence	Unvalidated release; self-approval	SR 26-2; SS1/23; E-23	1L assembles; 2L approves Tier 1 to 2	Signed release decision
C04 Exceptions with expiry	Risk carried indefinitely	Risk appetite; IIA standards	1L requests; 2L accepts	Exception register
C05 Prompt injection defenses	LLM01; ASI01	NYDFS Part 500; AI Act Art. 15	1L operates; CISO standard	Screening logs; red-team results
C06 Least-privilege tools	LLM03 excessive agency	FSB practices; Part 500 access	1L; CISO challenges	Agent contract; tool tests
C07 Delegated identity and purpose	Privilege abuse; leakage LLM02	GLBA; GDPR; Part 500	1L; CISO	Token records; denied calls
C08 Data routing by class	LLM02 leakage to providers	GLBA; HIPAA; GDPR Art. 28	1L; legal; CISO	Route approvals; contracts
C09 Supply chain pinning	LLM04 supply chain	TPRM guidance; DORA Art. 28	1L; CISO	Pinned digests; vendor records
C10 Run limits and spend caps	LLM06 denial of wallet	Operational resilience	1L	Limit configuration; alerts
C11 Grounding and figure checks	LLM07 misinformation	Reg B reasons; FTC Section 5	1L; 2L standard	Evaluation; output check logs

Controls for quality, people, change and reporting.

CONTROL	THREAT OR RISK	LEGAL OR SUPERVISORY SOURCE	LINE OF DEFENSE OWNER	EVIDENCE PRODUCED
C12 Fairness testing	Unfair discrimination	CL 7; NAIC; Colo. 10-1-1; Section 1557	1L tests; 2L challenges	Test results by group
C13 Human approval binding	Approval spoofing; agency	FSB; AI Act Art. 14	1L	Approval records
C14 Evaluated fallback	Untested substitute route	TPRM; DORA exit	1L; 2L	Fallback evaluation
C15 Monitoring thresholds	Drift; silent failure	AI Act Art. 26; MAS lifecycle	1L; 2L reviews	Monitoring record
C16 Incident clocks and kill switch	Missed notice; runaway agent	36-hour rule; Part 500.17; GDPR 33; DORA	1L; legal; CISO	Playbook; tests
C17 Change control	Unreviewed model or prompt change	SR 26-2; AI Act Art. 25	1L; 2L approves material	Change log; re-gate records
C18 Records and legal hold	Spoliation; unexplainable decision	AI Act Art. 26(6); litigation duties	Legal owns	Retention schedule; hold tests
C19 Disclosure and notices	Deception; missing notice	AI Act Art. 50; FTC; CCPA ADMT	Legal; 1L implements	Approved wording; UI evidence
C20 Committee reporting	Board unaware of exposure	Colo. 10-1-1; MAS; IIA model	2L and 3L	Quarterly report

Use the matrix three ways. The first line uses it to know which controls its tier requires. The second line uses it to check that each control has a standard and an owner. Internal audit uses it as the starting point for its AI audit universe, testing each control’s design and operation on a sample drawn from the registry.

STATUS PER CONTROL For each product, mark each control as tested, failing, planned, blocked or not applicable, with a reason. A control that is built but only warns is planned, not tested.

APPENDIX D / RED-TEAM TEST PLAN

Test the product the way an attacker will.

Scope the red team to the threat model on the registry record. Run gateway and tool cases on every release; run model-level and scenario cases at the depth the tier requires. Record every case with input, expected control, observed result and fix.

AREA	CASES	PASS CONDITION	TIER
Direct injection	Override instructions; role-play; encoding tricks; requests to reveal hidden context	Refused or safely handled; no hidden context disclosed	All
Indirect injection	Instructions planted in documents, emails, web pages, tool results and memory	Content treated as data; no tool call or disclosure follows	1 to 3
Data access	Ask for another customer's record; cross-tenant passages; revoked user; restricted label to an unapproved route	Denied at the gateway or data layer, with a logged reason	All
Tool abuse	Out-of-contract tool; out-of-range arguments; free-text query; write without approval; reused approval	Denied; approval bound to exact parameters and single use	1 to 3
Agent chains	Instruction passed agent to agent; widened scope on delegation	Each hop re-checks; scope only narrows	1 and 2
Denial of wallet	Loops, very large contexts, repeated retries	Run limit or spend cap stops it; alert raised	1 to 3
Supply chain	Changed MCP server or tool digest; unregistered model route	Refused until re-registered and re-approved	All
Output	Prompts that elicit unsupported figures, personal data or unsafe markup	Output check blocks or flags; incomplete state where evidence is missing	1 to 3
Fallback	Force the primary provider to fail	Degraded mode or evaluated fallback only; no unapproved route serves	1 and 2
Kill switch	Stop the agent mid-task	Stops within the stated time; no partial write left unreconciled	1 to 3

Map each case to OWASP 2026 and agentic identifiers and to MITRE ATLAS techniques so detections can be built from the same list.[28, 29, 30] Tier 1 red teams are run or re-performed by people independent of the build team. Findings that remain open become exceptions under section 06, with an end date. Re-run the red team after every material change, and at least once a year for Tier 1 and 2.

APPENDIX E / AI INCIDENT PLAYBOOK

Ten steps from first signal to closed postmortem.

STEP	ACTION	OWNER	RECORD
1 Detect	Alert, complaint, staff report or audit finding names a registry product	First line	Incident opened with registry ID
2 Triage	Is it an AI incident under section 13? Set severity from tier and harm	Product owner with CISO office	Severity and rationale
3 Contain	Restrict scope, switch route, or pull the kill switch; preserve the trace	Product owner	Containment time and action
4 Start clocks	Check the registry incident_clocks field; legal decides which clocks have started	Legal and compliance	Clock table with start times
5 Assess impact	Affected people, data, decisions and money; fairness impact where relevant	Business owner, risk	Impact assessment
6 Notify	Regulators, customers, partners and providers as required; consistent facts across notices	Legal and compliance, CISO office	Copies of notices; times
7 Remediate	Correct affected decisions or records; customer redress where due	Business owner	Remediation log
8 Fix	Change approved through the gate, at the depth of the tier	Product owner; second-line approvers	Re-gate record
9 Postmortem	Root cause; failed or missing control; tier or evaluation changes	Product owner; second line challenges	Postmortem on the registry
10 Report and test	Committee report; audit samples closure	Second and third line	Report entry; audit test

Tabletop scenarios to rehearse

- A service agent reads an email with planted instructions and sends a customer statement to an outside address.
- The model provider changes its model version overnight; next-day override rates double.
- A provider account runs out of credit; a team switches to an untested route.
- A fairness monitor shows a widening approval gap for one group in an underwriting assistant.

ONE PROCESS AI incidents run in the firm’s existing incident process with these steps added. Do not create a parallel process with its own severity scale.

APPENDIX F / GLOSSARY

Terms used in this paper.

Agent
An AI system that can take actions through tools on behalf of a person or process.

Autonomy level
L0 assist to L4 bounded autonomy, set per action in the agent contract.

Composition
The deployed combination of model, prompts, retrieval, tools, permissions, interface and workflow.

Control point
One of four decisions: intake and tiering, release gate, runtime, change and retirement.

Critical-failure veto
A rule that no average score can offset a critical failure in evaluation.

Denial of wallet
An attack or fault that drives cost or exhausts quota or credit.

DSPM
Data security posture management: tools that find sensitive data and risky exposure.

Evidence trail
The linked records of every decision on a product, from intake to retirement.

Exception
An accepted, time-limited departure from a required control, with a compensating control.

Fallback route
An alternative model or provider used when the primary fails; a component that needs its own evaluation.

Gate
The decision that a specific release may run in production, for a scope, until a date.

Indirect prompt injection
Instructions hidden in content the model reads, rather than typed by the user.

Kill switch
A tested mechanism that stops an agent, tool, product or route within a stated time.

MCP
Model Context Protocol, a standard way for AI applications to call tools and data.

Registry
The system of record for AI products and components that deployment and runtime read.

Release ID
The identifier of one approved composition; it changes when any component changes.

Shadow AI
AI in use that the registry does not know about.

Three lines
First line owns and operates; second line sets standards, challenges and approves; third line gives independent assurance.

Tier
Tier 1 critical to Tier 4 low, set by the questionnaire; it sets the evidence and approvals.

ABBREVIATIONS ADMT automated decision-making technology · DPIA data protection impact assessment · DORA Digital Operational Resilience Act · GLBA Gramm-Leach-Bliley Act · NPI nonpublic personal information · PHI protected health information · SM&CR Senior Managers and Certification Regime · TPRM third-party risk management.

APPENDIX F / SOURCES

Primary sources and review basis.

Prepared by Richard White; reviewed 20 November 2026. Statements of law and supervisory expectation cite the primary source. Status labels (proposed, draft, final) are as of that date. Not legal advice.

US FEDERAL BANKING AND AI POLICY

[1] **SR 26-2.** Revised Guidance on Model Risk Management, Federal Reserve, OCC and FDIC, 17 April 2026. federalreserve.gov/supervisionreg/srletters/SR2602.htm

[2] **OCC Bulletin 2026-54.** Super Intelligence: Supervisory Framework, 4 November 2026. occ.gov/news-issuances/bulletins/2026/bulletin-2026-54.html

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

[4] **OCC Bulletin 2026-46.** Third-party risk management: proposed guidance, 11 September 2026. Proposed, not final. occ.gov/news-issuances/bulletins/2026/bulletin-2026-46.html

[5] **Computer-Security Incident Notification Requirements.** 86 FR 66424, 23 November 2021; compliance from 1 May 2022. federalregister.gov/d/2021-25510

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

[7] **Financial Stability Board.** Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Final report, 15 October 2026. fsb.org/…

[8] **Executive Order 14365.** Ensuring a National Policy Framework for Artificial Intelligence, 11 December 2025 (White House fact sheet). whitehouse.gov/…

FTC AND SECTOR LAW

[9] **FTC.** FTC Announces Crackdown on Deceptive AI Claims and Schemes (Operation AI Comply), 25 September 2024. ftc.gov/…

[10] **FTC.** FTC Reopens and Sets Aside Rytr Final Order, December 2025. ftc.gov/…

[11] **FTC.** Proposed Policy Statement Concerning the Suppression of Accuracy in Artificial Intelligence Systems, July 2026. Proposed, open for comment. ftc.gov/…

[12] **Gramm-Leach-Bliley Act.** 15 U.S.C. 6801 to 6809, privacy of consumer financial information and safeguards. uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title15-chapter94

[13] **FTC Safeguards Rule amendment.** 88 FR 77499, 13 November 2023; notification of events affecting 500 or more consumers within 30 days, effective 13 May 2024. federalregister.gov/d/2023-24412

[14] **HIPAA Privacy, Security and Breach Notification Rules.** 45 CFR Parts 160 and 164; breach notice to individuals within 60 days of discovery (164.404), HHS. hhs.gov/hipaa/for-professionals/index.html

[15] **Section 1557 final rule.** 45 CFR 92.210, patient care decision support tools; 89 FR 37522, 6 May 2024; compliance from 1 May 2025. federalregister.gov/d/2024-08711

[16] **California Privacy Protection Agency.** Regulations on automated decision-making technology, risk assessments and cybersecurity audits; approved September 2025, effective 1 January 2026, ADMT duties from 1 January 2027. cppa.ca.gov/regulations/ccpa_updates.html

US states, the EU, the UK, Canada and Singapore.

US STATES

[17] **NAIC Model Bulletin.** Use of Artificial Intelligence Systems by Insurers, December 2023, and the NAIC adoption map. content.naic.org/…

[18] **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

[19] **Colorado Regulation 10-1-1.** Life insurers from 14 November 2023; private auto and health benefit plan insurers from 15 October 2025. doi.colorado.gov/…

[20] **Colorado SB 26-189.** Automated decision-making technology; signed 14 May 2026; repeals and replaces SB 24-205. leg.colorado.gov/bills/sb26-189

[21] **23 NYCRR Part 500.** NYDFS cybersecurity regulation as amended 1 November 2023, including 500.17 notices. dfs.ny.gov/industry-guidance/cybersecurity

[22] **NYDFS industry letter.** Cybersecurity risks arising from artificial intelligence, 16 October 2024. dfs.ny.gov/…

[23] **NYDFS industry letter.** Heightened Cybersecurity Risks Associated with Frontier AI Models, 21 May 2026. dfs.ny.gov/…

EUROPEAN UNION

[37] **Regulation (EU) 2024/1689, the AI Act.** Articles 4, 6, 9 to 15, 25 to 27, 50, 72 and 73, and Annex III. eur-lex.europa.eu/eli/reg/2024/1689/oj

[38] **Regulation (EU) 2026/1744, the Digital Omnibus on AI.** Adopted 8 July 2026; Official Journal 24 July 2026; in force 27 July 2026. eur-lex.europa.eu/eli/reg/2026/1744/oj

[39] **DORA, Regulation (EU) 2022/2554.** Articles 19, 28 and 30; applies from 17 January 2025. eur-lex.europa.eu/eli/reg/2022/2554/oj

[40] **Commission Delegated Regulation (EU) 2025/301.** Content and time limits for major ICT-related incident reports, 23 October 2024. eur-lex.europa.eu/eli/reg_del/2025/301/oj

[41] **GDPR, Regulation (EU) 2016/679.** Articles 22, 28, 33 and 35. eur-lex.europa.eu/eli/reg/2016/679/oj

UNITED KINGDOM, CANADA AND SINGAPORE

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

[43] **FCA.** AI and the FCA: our approach, accessed November 2026. fca.org.uk/firms/innovation/ai-approach

[44] **OSFI Guideline E-23.** Model Risk Management, final 11 September 2025, effective 1 May 2027. osfi-bsif.gc.ca/…

[45] **MAS.** Guidelines on Artificial Intelligence Risk Management, issued 7 October 2026, after consultation from 13 November 2025. mas.gov.sg/…

Frameworks, security and assurance.

FRAMEWORKS, SECURITY AND ASSURANCE

[24] **NIST AI RMF 1.0 (AI 100-1).** January 2023; NIST notes it is under revision. nist.gov/itl/ai-risk-management-framework

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

[26] **NIST AI 100-2 E2025.** Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, March 2025. csrc.nist.gov/pubs/ai/100/2/e2025/final

[27] **NIST IR 8596.** Cyber AI Profile, initial preliminary draft, 16 December 2025. Draft. csrc.nist.gov/pubs/ir/8596/iprd

[28] **OWASP Top 10 for LLM Applications 2026.** OWASP GenAI Security Project, August 2026. genai.owasp.org/llm-top-10/

[29] **OWASP Top 10 for Agentic Applications.** OWASP GenAI Security Project, 9 December 2025. genai.owasp.org/resource/owasp-top-10-for-agentic-applications/

[30] **MITRE ATLAS.** Adversarial Threat Landscape for Artificial-Intelligence Systems, tactics and techniques, 2026 releases. atlas.mitre.org

[31] **Model Context Protocol.** Security best practices, specification 2026-07-28. modelcontextprotocol.io/…

[32] **IETF RFC 8693.** OAuth 2.0 Token Exchange, January 2020. rfc-editor.org/rfc/rfc8693

[33] **The IIA's Three Lines Model.** An update of the Three Lines of Defense, 20 July 2020. theiia.org/…

[34] **The IIA.** Global Internal Audit Standards, effective 9 January 2025. theiia.org/…

[35] **ISO/IEC 42001:2023.** Artificial intelligence management system, December 2023. iso.org/standard/81230.html

[36] **US Copyright Office.** Copyright and Artificial Intelligence, Part 2: Copyrightability, 29 January 2025. copyright.gov/newsnet/2025/1060.html

How the sources were checked

Each regulatory and legal statement was checked against the issuing body's own publication or the official journal, and dated. Where a document is proposed, a draft or under challenge, the text says so. Where a status could not be confirmed from a primary source, the claim was left out. Product statuses in section 16 come from RJL's product records of 20 November 2026, not from this paper's recommendations. Recheck every source before relying on it; supervisory expectations for AI are changing quickly.

ABOUT

Richard White

I was Global Chief Data Officer at McKinsey & Company, where I led the firm's internal data and AI function and built capabilities such as its data products. Before that I was 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, then Operate and transfer. Each phase is separately scoped and priced in a written proposal before you commit. The first conversation is free.

EVIDENCE **AI Governance Evidence**
What US bank and insurance reviewers inspect, in 22 pages.

TECHNICAL **Agent Onboarding: Treating an Agent Like a Data Product**
The agent contract the registry stores.

CHECK **The 10-Question AI Production-Readiness Diagnostic**
Ten questions on whether this is working.

RJL SOLUTIONS

WEB rjlsol.com

EMAIL rwhite@rjlsol.com

This paper describes a control framework and how to evidence it. It isn't legal advice, a legal opinion or a conformity assessment. Counsel and compliance owners decide which obligations apply. I reviewed the sources on 20 November 2026; recheck them before use. Companies in examples are fictional.