RISK & COMPLIANCE MAPPING

Six frameworks. One operating model.

Six frameworks, one operating model. The EU AI Act, ISO/IEC 42001, NIST AI RMF, SR 11-7, PRA SS1/23 and DORA each ask for different evidence; produce it six separate ways and compliance becomes the bottleneck. This maps every clause onto the four layers, so the team produces the evidence once and orients it for each reader: the cross-framework matrix, the five-page conformity report, and the cadence each framework expects. For Heads of Risk, Compliance, Procurement and vendor risk.

FOR

Head of Risk, Compliance, Procurement, vendor risk

FRAMEWORKS

AI Act, ISO 42001, NIST AI RMF, SR 11-7, PRA SS1/23, DORA

COMPANION

CISO & vendor briefing for security posture

DIRECT LINE

charafeddine@cohorte.co

02 THE PROBLEM WITH SIX FRAMEWORKS

A regulator wants one document, not five.

Most enterprises operating AI in 2026 are subject to three to five concurrent frameworks: the EU AI Act, ISO/IEC 42001 if certified, NIST AI RMF if US-aligned, SR 11-7 or PRA SS1/23 if regulated banks, DORA if EU financial services. Each framework is rigorous on its own. Stacked, they generate redundant evidence requests, conflicting cadences, and a stack of compliance dashboards that no single owner reads.

The mapping in this briefing solves one problem: the evidence the four-layer operating model produces (the scoping briefs, the reliability levels, the agent passports, the monitoring logs) is the same evidence each framework asks for, presented in the language each framework uses. The team installs the operating model once. The mapping reorients the artefacts toward whichever framework is asking.

The orientation matters. A reliability level satisfies AI Act Article 15 (accuracy, robustness), ISO 42001 clause 8.4 (performance evaluation), NIST AI RMF Measure-2.5 (system performance), SR 11-7 model-monitoring requirements, and PRA SS1/23 paragraph 4.20 simultaneously. They are five names for the same evidence. The mapping makes this visible so the team produces the evidence once and the compliance function annotates it five ways.

The gaps matter too. Not every framework asks for everything. DORA has specific requirements on third-party risk that the others handle elsewhere. SR 11-7 has model-risk-tiering requirements that have no direct AI Act analogue. The mapping makes those framework-specific obligations visible so the team does not assume the operating model covers them. Where it does not, the briefing names the gap and what to add.

> Comply with six frameworks. Build the evidence once.

03 EU AI ACT

EU Regulation 2024/1689.

High-risk obligations and their evidence.

The AI Act enters full enforcement for high-risk systems in August 2026. The articles below are the ones that require operational evidence. Each maps onto a layer of the Cohorte stack and onto an artefact the operating model already produces.

REFERENCE	WHAT IT REQUIRES	WHERE IT LIVES IN THE OPERATING MODEL
Art. 6 + Annex III	High-risk classification of the system based on use case.	Layer 01 (scoping). The brief declares the Annex III category. The no-list captures systems classified out of high-risk.
Art. 9	Risk management system across the AI lifecycle.	The four-layer stack is the lifecycle. Documented across all four briefings. **This briefing maps the artefacts to the article's requirements.**
Art. 10	Data governance: relevant, representative, error-free training and test data.	Layer 02 (verification). Calibration set governance: provenance, labelling rules, versioning, exchangeability check.
Art. 11	Technical documentation, before placing on the market.	Agent passport + verification artefact + scoping brief. **The registry assembles the documentation on demand.**
Art. 12	Automatic logging of events over the system's lifetime.	Layer 04 (monitoring). Structured logs with selective full-trace retention. agent-monitor reference repo.
Art. 13	Transparency: information to deployers about the system's capabilities and limits.	The agent passport is the document. Reliability level, abstention behaviour, scope of authority all included.
Art. 14	Effective human oversight by natural persons.	Layer 03 (governance). Mind-in-the-loop gate placement, escalation paths, named owner.
Art. 15	Accuracy, robustness and cybersecurity, with metrics and measurement.	Layer 02 (verification): reliability level. Layer 04: drift signals. Briefing 03: exploitation-surface countermeasures.
Art. 16, 26	Provider and deployer obligations during operation.	Registry change log. Incident grid. Post-deployment monitoring evidence.

04 ISO/IEC 42001

ISO The AI management system.

Clauses 4 to 10, mapped.

ISO/IEC 42001:2023 is the management-system standard for AI. It follows the structure of ISO 9001 and ISO 27001: a Plan-Do-Check-Act cycle around context, leadership, planning, support, operation, performance evaluation, and improvement. The clauses below are the ones that materially shape the operating model.

CLAUSE	WHAT IT REQUIRES	WHERE IT LIVES IN THE OPERATING MODEL
4: Context	Internal and external context, interested parties, AI scope.	The scoping brief plus the no-list. The systems in scope and the explicitly excluded.
5: Leadership	AI policy, roles and responsibilities, accountability.	The agent passport names the owner. The registry assigns accountability per system, not per committee.
6.1: Risk	Risks and opportunities, AI risk assessment, treatment.	Layer 01 produces the no-list (refused risks). Layer 03 produces the residual risk in the passport. Layer 04 detects emerging risk.
6.2: Objectives	AI objectives, plans to achieve them.	The reliability bar declared in the scoping brief. The retirement condition. The cadence of re-calibration.
7: Support	Resources, competence, awareness, communication, documented information.	Cohorte programs. The bootcamp installs competence. The Curriculum License sustains it. The registry holds the documented information.
8.1: Op. planning	Operational planning and control of AI activities.	The four-layer stack as a loop. The cadence definitions from Briefing 02 (calibration) and Briefing 03 (governance).
8.2–8.5	AI system impact, lifecycle, data, third-party.	Layer 01 (impact), all four layers (lifecycle), Layer 02 (data), CISO briefing + Briefing 05 (third-party / sovereignty).
9: Performance eval.	Monitoring, measurement, analysis, evaluation, internal audit, management review.	Layer 04. Monitoring signals. Weekly governance report. Quarterly external calibration. The conformity report.
10: Improvement	Continual improvement, non-conformity and corrective action.	The incident grid. Post-mortem rolling into passport change log. The next quarter's calibration set incorporates the lessons.

05 NIST AI RMF

NIST Govern, Map, Measure, Manage.

Plus the GenAI Profile.

The NIST AI Risk Management Framework (AI RMF 1.0) is the dominant US-aligned voluntary framework. The four functions decompose AI risk management into Govern (cross-cutting culture), Map (context), Measure (analysis), and Manage (action). The 2024 Generative AI Profile adds GenAI-specific actions on top.

FUNCTION	REPRESENTATIVE CATEGORIES	WHERE IT LIVES IN THE OPERATING MODEL
Govern	Govern-1 (policies, processes). Govern-2 (accountability structures). Govern-3 (workforce). Govern-4 (culture). Govern-5 (engagement). Govern-6 (third-party).	The registry (accountability). The passport (responsibilities). Cohorte programs (workforce + culture). The CISO briefing (third-party).
Map	Map-1 (context established). Map-2 (categorisation). Map-3 (capabilities, risks). Map-4 (impact). Map-5 (stakeholder).	Layer 01 (scoping brief, no-list). The Annex III tagging. The impact section of the brief signed by the executive sponsor.
Measure	Measure-1 (identified methods). Measure-2 (evaluated). Measure-2.5 (system performance and reliability). Measure-3 (mechanisms tracked). Measure-4 (feedback integrated).	Layer 02 (the reliability level). Layer 04 (the five monitoring signals). The calibration cadence.
Manage	Manage-1 (risks prioritised). Manage-2 (treated). Manage-3 (third-party). Manage-4 (post-deployment).	The incident grid. The hold action. The retirement condition. The post-mortem feeding the next calibration.
GenAI Profile	CBRN, confabulation, dangerous content, data privacy, environmental, human-AI configuration, information integrity, information security, intellectual property, obscene content, toxicity, value chain.	The exploitation-surface paper covers information security (Briefing 03). The CISO briefing covers data privacy. The scoping no-list refuses categories the operating model is not equipped to govern.

> The four functions are the structure. The Profile is the GenAI-specific delta. Both map onto the same evidence.

06 SR 11-7 + PRA SS1/23

MR Model risk.

The framework banks audit against.

For banks operating in the US (SR 11-7), the UK (PRA SS1/23), and most G20 jurisdictions following similar guidance, AI systems fall under the model-risk-management framework. The framework is older than LLMs and more rigorous than the AI-specific frameworks. The mapping below shows where the operating model satisfies it and where additional model-risk-specific artefacts are required.

REQUIREMENT	WHAT IT ASKS FOR	WHERE IT LIVES
SR 11–7 §III	Model development, including conceptual soundness, ongoing monitoring, outcomes analysis.	Layer 01 (scoping brief documents the conceptual basis). Layer 02 (verification). Layer 04 (outcomes analysis through monitoring).
SR 11–7 §IV	Model validation, including evaluation by independent parties, ongoing monitoring, outcomes analysis.	Independent validation requires a second-line review function. **This is an addition to the operating model, not a substitute.** The artefacts (passport, reliability level, calibration set) are the inputs.
SR 11–7 §V	Governance, policies, controls. Model inventory.	The registry is the inventory. The governance committee is the controls owner. The passport is the per-model record.
PRA SS1/23 §4	Governance, model definition, risk management, lifecycle, accountability, data, model development, model validation, model deployment, monitoring.	Maps onto all four layers. The PRA's "model definition" lines up with the scoping brief. "Accountability" lines up with the named owner in the passport.
PRA SS1/23 §4.20	Continuous monitoring of model performance.	Layer 04 (the five signals). The continuous calibration cadence from Briefing 02.

The model-risk-management addition. SR 11-7 and PRA SS1/23 require independent second-line validation of the model. The operating model produces the artefacts the second-line team needs, but the second-line function itself is a regulated organisational structure that the operating model does not provide. A bank using Cohorte programs will have its independent validation team review the artefacts; the team itself remains the bank's organisational responsibility.

07 DORA

Operational resilience.

Third-party AI as critical ICT service.

The EU Digital Operational Resilience Act (Regulation 2022/2554) applies to EU financial entities. AI vendors offering critical services come under its third-party risk regime. The operating model produces evidence DORA needs; this section names the DORA-specific obligations the operating model does not, by itself, satisfy.

ARTICLE	WHAT IT REQUIRES	WHERE IT LIVES IN THE OPERATING MODEL
Art. 5–14	ICT risk management framework. Identification, protection, detection, response, recovery.	The four-layer stack is the ICT risk management framework for the AI portfolio. **Detection and response are covered by Layer 04 and the incident grid.**
Art. 15–16	ICT incident reporting. Classification, root-cause analysis, lessons learnt.	The incident grid (Briefing 03). The post-mortem. The change log. **The reporting cadence to ECB/EIOPA is an addition; the operating model produces the inputs.**
Art. 24–27	Digital operational resilience testing, including threat-led penetration testing.	The exploitation-surface assessment (Briefing 03) is the AI-specific addition. Conventional pen-testing remains required and unchanged.
Art. 28–30	Third-party ICT risk management. Register of contractual arrangements.	**This is the DORA obligation the operating model most clearly extends.** The agent registry tracks Cohorte-installed systems; the DORA register tracks the contractual relationships with the underlying vendors (model providers, cloud, retrieval systems).
Art. 31–44	Oversight of critical third-party ICT service providers.	If the AI service is provided by a third-party meeting DORA's criticality criteria, the entity sits under direct supervision. The operating model surfaces dependency to support the criticality assessment.

> DORA is the framework that asks who the vendors are. The registry is the document that answers.

08 THE CROSS-FRAMEWORK MATRIX

One artefact. Six frameworks. The mapping at a glance.

The matrix below is the page a Head of Compliance can hand to procurement and to second-line audit. It says: this is the artefact we produce, and these are the framework citations for which the artefact is the evidence. The team produces the artefact once. The matrix annotates it for each audience.

ARTEFACT	AI ACT	ISO 42001	NIST AI RMF	SR 11-7	DORA
Scoping brief + no-list	Art. 6, 9 (Annex III, risk mgmt)	4, 6.1 (context, risk)	Map-1, Map-2 (context)	§III conceptual	Art. 5-14 ICT risk
Agent passport	Art. 11, 13 (documentation, transparency)	5, 7.5 (leadership, doc info)	Govern-2 (accountability)	§V inventory	Art. 28-30 third-party register input
Reliability level + cal. set	Art. 15 (accuracy, robustness)	8.4, 9.1 (perf eval)	Measure-2, 2.5	§III conceptual + ongoing monitoring	Art. 5-14 detection
Gate + escalation	Art. 14 (human oversight)	8.1 (operational control)	Manage-2 (treated)	§IV.B model validation	Art. 5-14 response
Monitoring log	Art. 12 (logging)	9.1 (monitoring)	Measure-3 (tracked)	§III ongoing monitoring	Art. 5-14 detection + Art. 15-16 incident
Incident grid + post-mortem	Art. 26 (deployer obligations)	10 (improvement)	Manage-4 (post-deployment)	§V controls	Art. 15-16 incident reporting
Calibration cadence	Art. 15 (over the lifecycle)	10 (continual improvement)	Measure-4 (feedback)	§III.D ongoing	Art. 24-27 testing
Exploitation-surface assessment	Art. 15 (cybersecurity)	8.5 (third-party)	GenAI Profile (info security)	§IV.B validation	Art. 24-27 threat-led

The matrix is not a substitute for the framework. A regulator will read the framework, not the matrix. The matrix exists so the team can produce one set of artefacts and orient them five ways for five readers. The framework-specific obligations the matrix does not cover (independent validation in SR 11-7, third-party register in DORA, AI policy in ISO 42001 clause 5) require additional artefacts produced outside the operating model.

09 THE CONFORMITY REPORT

What the team hands the regulator. Five pages.

The conformity report is the document a deployer of a high-risk AI system produces under the AI Act, by extension under ISO 42001 internal audit, by re-orientation under NIST AI RMF reporting, and as part of the SR 11-7 effective challenge package. Five pages. Generated from the registry. The format below is what Cohorte teaches teams to produce, and what the AI Readiness Program installs.

1 **System identity & classification**
Name, registered purpose, owner, date placed on the market, Annex III category (or a stated reason it is not high-risk), model and version, the integration architecture in one diagram.
ANSWERS: AI ACT ART. 11

2 **Risk management & verification**
The risk assessment from the scoping brief, the reliability level, the calibration set and date, the deviation history over the last four calibrations, the retirement condition.
ANSWERS: AI ACT ART. 9 & 15 · SR 11-7 VALIDATION

3 **Human oversight & authority**
Gate placement, the reviewer's role and training, the escalation path, the agent's permission profile and its change history.
ANSWERS: AI ACT ART. 14

4 **Monitoring & incidents**
The five signals across the reporting window; the incident log with severity, root cause and remediation, and the effect of each fix on the reliability profile.
ANSWERS: AI ACT ART. 12 & 72 · DORA INCIDENT REPORTING

5 **Lifecycle & disposal**
The change log and the reasoning behind any material change; the retirement plan, or the disposal report and data-handling closure once retirement has happened.
ANSWERS: ISO/IEC 42001 · AI ACT ART. 17

> A conformity report that takes five pages because the system was designed to produce one is the difference between governance and theatre.

10 CADENCE BY FRAMEWORK

When each framework expects a report. And which one fires first.

The frameworks have different reporting calendars. The team that installs them once but reports on the framework-specific cadence will find that the deepest reporting obligation tends to drive the cadence of the others. The schedule below is what Cohorte teaches teams to track.

FRAMEWORK	TRIGGER	CADENCE
AI Act	Placement on the market; substantial modification; serious incident.	Conformity assessment at placement; technical documentation maintained; **serious incidents within 15 days** (Art. 73).
ISO 42001	Certification audit cycle.	Internal audit at planned intervals; **management review at planned intervals**; certification audit annually.
NIST AI RMF	Voluntary; cadence set by the adopting organisation.	Cohorte recommends **quarterly Measure reporting + annual Govern review** aligned with ISO 42001 cycle.
SR 11–7	Continuous; tiered by model materiality.	Annual model inventory; **ongoing monitoring on a frequency commensurate with materiality**; periodic validation.
PRA SS1/23	Continuous; firm responsible for proportional cadence.	Continuous monitoring per §4.20. **Material model validation at minimum annually**; senior managers report at appropriate intervals.
DORA	Continuous; major ICT incident triggers report.	Major ICT incident initial notification **within 4 hours**; intermediate report within 72 hours; final report within one month.

The fast cadences set the engineering bar. A team that builds for the 4-hour DORA notification builds for everything else. A team that builds for the annual ISO 42001 internal audit will fail the DORA test. The operating model defaults to the fast cadence; the framework with the slowest cadence consumes the artefacts the fast cadence produces.

11 WHAT THIS IS NOT

Three claims this briefing refuses to make.

A compliance-mapping document that does not name its limits is one the legal team will not trust. The three statements below are deliberately on the page.

Not legal advice. The mappings on the previous pages reflect a competent practitioner's reading of the frameworks as published. They are not legal opinions. A regulated entity will validate the mapping with its general counsel and, for the AI Act high-risk classification, with a notified body. A firm that ships against this mapping without legal review treats it as more than it claims to be.

Not a substitute for the framework. A regulator reads the framework. The matrix on page 08 is a pedagogical aid; it is not the framework. Where the matrix is silent, the framework is not. The four-layer operating model satisfies a large fraction of the frameworks' operational requirements; it does not satisfy the structural requirements (independent validation, AI policy ratification, board-level reporting cadence) that belong to the organisation, not the operating model.

Not a substitute for the certification programme. ISO/IEC 42001 certification requires an external audit by an accredited body. Cohorte programs install the management system; they do not award the certification. A firm that intends to certify uses Cohorte to be ready and a different organisation to be audited.

> A briefing that respects the regulator and the buyer is one that says what it does not do. This one does.

12 HOW THIS LANDS

The AI Readiness Program. Where the mapping is installed.

The mapping in this briefing is the operating content of the AI Readiness Program, the deepest Cohorte engagement. The Team Bootcamp installs the operating model on one team's systems; the License rolls out the curriculum across the firm; the AI Readiness Program installs the operating model *and* wires it to the regulatory reporting cadence the firm is subject to.

MONTH 1	**Framework triage.** Which frameworks apply, with which intensity. The mapping matrix tailored to the firm's regulatory perimeter.
MONTH 2	**Inventory and gap.** The current portfolio audited against the matrix. Gaps named in writing with the framework citations behind them.
MONTH 3	**Operating-model installation.** The four layers wired across the portfolio. Registry populated. Passports drafted for high-risk slice first.
MONTH 4-5	**Cadence installation.** Reporting calendars wired to the frameworks' cadences. The fast cadence (DORA, AI Act incident) drives the architecture.
MONTH 6	**First conformity report.** Generated from the registry. Reviewed by the firm's second-line risk function. Adjusted. Defended.

The Readiness Program is co-delivered with the second-line risk function. Cohorte does not replace the firm's compliance team. Cohorte teaches the operating model; the compliance team owns the framework interpretation. The artefacts produced are the joint output of the engagement and the compliance function. The deliverable is the conformity report that both sign.

13 FROM THE FIELD

The PwC AI Factory baseline. Built before enforcement.

The conformity-reporting baseline for the systems running in the PwC France & Maghreb AI Factory was built between 2024 and 2026, before the EU AI Act enforcement date for high-risk systems. The reference is Patrick Monteiro, CIO; the call is arranged after a mutual NDA.

FRAMEWORK COVERAGE	**EU AI Act** (Articles 9-17 mapped). **ISO/IEC 42001** (clauses 4-10 mapped). **NIST AI RMF** Govern + Measure functions. The other frameworks are layered as the firm's portfolio extends to them.
CONFORMITY-REPORTING BASELINE	The five-page conformity report format from page 09 was developed inside this engagement. **The same format is used across high-risk systems** and is the audit-facing artefact.
CADENCE	Weekly internal reliability cycle. Monthly cross-portfolio review. Quarterly external calibration. **Annual management review with the named second-line risk function.**
OUTCOME	A portfolio of 60+ production systems on the registry, each with the operating-model evidence, available to the regulator on request. **No findings on the first audit cycle in the registry layer.**

> A compliance baseline built before the enforcement date reads differently in audit than one assembled in the quarter after.

14 REFERENCES

References. Each framework, in source.

Reading list. Where to find the source text for each framework cited on the previous pages.

European Union (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L series.

European Commission (2024-2026). Guidelines on the application of the AI Act. Ongoing publication. The Commission's AI Office maintains the canonical interpretation.

ISO/IEC (2023). ISO/IEC 42001:2023 Information technology, Artificial intelligence, Management system. Geneva: ISO.

ISO/IEC (2023). ISO/IEC 23894:2023 Information technology, Artificial intelligence, Guidance on risk management. Geneva: ISO. Companion document to 42001 on AI risk specifically.

NIST (2023). AI Risk Management Framework (AI RMF 1.0). NIST AI 100-1. National Institute of Standards and Technology, US Department of Commerce.

NIST (2024). AI RMF Generative AI Profile. NIST AI 600-1. The GenAI-specific actions overlaid on the 100-1 functions.

Federal Reserve (2011, restated). SR 11-7 Supervisory Letter on Guidance on Model Risk Management. Board of Governors of the Federal Reserve System.

Prudential Regulation Authority (2023). SS1/23 Model risk management principles for banks. Bank of England.

European Union (2022). Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector (DORA).

European Supervisory Authorities (2024-2026). DORA Regulatory Technical Standards (RTS) and Implementing Technical Standards (ITS). The detail behind the regulation.

OWASP (2025). OWASP Top 10 for Large Language Model Applications, 2025 release.

Mouzouni, C. (2026). The three Cohorte research papers and the trust-and-governance briefings 01-05. Full record at **teams.cohorte.co/research**.

FOR HEADS OF RISK AND COMPLIANCE

One discovery call with your compliance lead.

Sixty minutes. Bring the framework perimeter. We walk the mapping against your portfolio, name the gaps, and you leave with the cross-framework matrix tailored to the firm.

charafeddine@cohorte.co

Cohorte SAS · Société par actions simplifiée, registered in France · founded September 2022 · Paris & Rabat

The full mapping work product, the framework reading list, and the conformity-report templates · at **teams.cohorte.co/research**