CXO INTELLIGENCE / GOVERNANCE SERIES / EDITION ONE / 2026

The Enterprise AI Governance Guide and Control Checklist

Converging the EU AI Act, the NIST AI Risk Management Framework, ISO/IEC 42001 and Singapore's Model AI Governance Framework into one operational system

INCLUDES THE 27 SHEET COMPANION WORKBOOK, FREE AT AKHAWAT.COM/TOOLKITS

21	105	19	15
CHAPTERS	CONTROLS	JURISDICTIONS	POLICIES

Prashant Akhawat • Building AI-Native Enterprises • akhawat.com

Version 1.0 • Information current as of August 2026 • Not legal advice

COMPANION WORKBOOK, INCLUDED FREE

The Enterprise AI Governance Toolkit

This guide tells you what to do. The workbook is where you do it. One Excel file, 27 connected sheets, no macros and no sign-up wall inside the file.

What it does that a template does not

- **It classifies for you.** Answer ten questions about a system and the workbook derives the risk tier, the compliance deadline, the control load, and whether you have become a GPAI provider without deciding to.
- **It flags the regimes that bind you.** Mark your markets and 19 jurisdictions self-flag as applies, monitor or out of scope, each mapped back to the same 32 capabilities.
- **The gap report writes itself.** Every open P1 foundation, in priority order, updating as you work. Pick a name and see only that person's outstanding items.
- **It tells you what to do next** in plain English on the Command Center, and the answer changes as you progress.
- **It carries the registers most programs lack:** risk, incidents, vendors, impact assessments, model cards, training records, statement of applicability, and a RACI that flags any activity without exactly one accountable owner.

Also inside

- A 260 action implementation checklist across 17 domains, prioritized P1 to P3
- The 32 capability convergence crosswalk, and a clause level map of the wider ISO AI family
- A ten dimension maturity assessment with a radar profile and the gap to Level 3
- A 21 metric board pack, and a glossary of 115 terms so the file stands alone

DOWNLOAD, FREE

akhawat.com/toolkits

Current release of the guide and the workbook. Versioned, not archived.

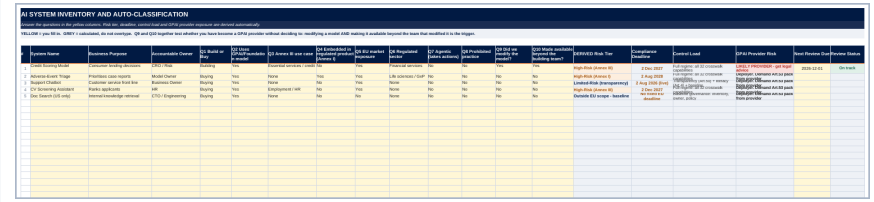

AI SYSTEMS SHEET. The three grey columns are derived, not typed.

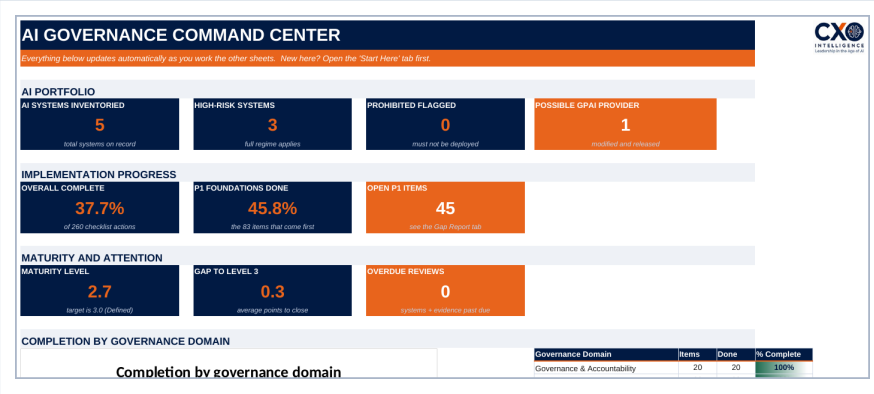

COMMAND CENTER. Updates as you work the other sheets.

CONTENTS

What is in this guide

How this guide was validated

Every regulatory position in this document was checked against primary or official sources in August 2026 rather than reproduced from earlier summaries. That verification changed two entries materially.

Colorado repealed and replaced its 2024 AI Act. The duty of care and the deployer impact assessment obligations that most published guidance still describes were removed. **Canada's** Artificial Intelligence and Data Act lapsed with Bill C-27 in January 2025 and has not been reintroduced, so any roadmap containing an AIDA row is planning against a statute that does not exist.

References to NIST AI RMF are made at **function level** (GOVERN, MAP, MEASURE, MANAGE), which is verified against NIST AI 100-1. Subcategory identifiers appear only where individually confirmed. This is deliberate: a wrong subcategory number in a document you execute from is worse than a correct function reference.

Standards are referenced by number, clause and intent only. **No normative or copyrighted standard text is reproduced anywhere in this document.** Implementing against any ISO/IEC standard requires a licensed copy.

> **A NOTE ON CURRENCY**
>
> AI regulation moves faster than any static document. This guide states its verification date openly and is versioned rather than assigned a permanent identifier, precisely because it will need replacing. Re-verify against primary sources before making a compliance decision, and check akhawat.com/toolkits for the current release.

CHAPTER 1

How to use this document

An execution document. It states what to do, who is accountable, what artifact proves it was done, and which provision it satisfies.

What this is

It assumes you have already accepted that AI governance is necessary and are now responsible for producing a defensible program. It is not legal advice, not a substitute for the primary sources, and not a certification. Where the law is unsettled, it says so rather than resolving the ambiguity for you.

What is where

If you need	Go to
To know whether any of this applies to you	Chapter 2
To know what you must stop doing immediately	Chapter 3
The single control set behind everything else	Chapter 5
The controls themselves, by domain	Chapters 6 to 14
Something to present	Chapter 15, marked liftable
Where to start on Monday	Chapter 16
What to write down	Chapter 17

If you need	Go to
Whether you are ready to certify	Chapter 18
What you owe if you build on a foundation model	Chapter 19
How to lawfully ship a high risk system	Chapter 20
What applies outside the EU, and in your sector	Chapter 21

Four reading paths

Reader	Read	Skim	Skip
Board and audit committee	Ch 2 scope summary, Ch 3 in full, Ch 15 board pack	Ch 4, Ch 16	Control tables
Legal and compliance	Ch 2, Ch 3, Ch 11, Ch 14 reporting, Ch 17	Ch 5, Ch 9	Ch 12
Risk and audit	Ch 2, Ch 7, Ch 9, Ch 15, Ch 18	Ch 5, Ch 10, Ch 14	Ch 12 detail
Engineering and security	Ch 9, Ch 12, Ch 13, Ch 14	Ch 6, Ch 8, Ch 10	Ch 4, Ch 17

Every reader should read Chapter 2 and Chapter 3. Scope determines whether the rest applies, and the red lines are the only place where the correct response is to stop rather than to control.

Notation

Control identifiers. Three-letter domain prefix plus sequence, for example EVD-002. Identifiers are stable across versions. Retired controls are not reused.

Applies when. Compressed as role / tier / jurisdiction.

Code	Meaning
P, D, I, X	Provider, Deployer, Importer or distributor, Product manufacturer
PR, HR, LR, MR	Prohibited, High risk, Limited risk, Minimal risk
GP	General purpose AI model track
EU, US-CA, SG, IN, ALL	Jurisdiction, or ALL where the control is good practice everywhere

P/HR/EU reads: applies to providers of high risk systems with an EU nexus.

Validity condition. Every control names the change events that void its evidence. A control with a satisfied validity condition is not satisfied. This is deliberate and is explained in Chapter 9.

Verification markers. A reference marked with a dagger was not verifiable against a primary source at the time of writing and appears in the verify list at Appendix C. Treat it as indicative until confirmed.

Three design principles carried throughout

These are stated once here and applied without further argument.

1. Evidence is emitted by the enforcement point, not narrated around it.

A log written by a reporting layer records what was reported. A log written by the component that authorizes the action records what happened. Where these differ, only the second is evidence.

2. Completeness is a separate failure mode from tamper evidence, and is settled at the write path.

Cryptographic integrity proves that the records you hold were not altered. It says nothing about the records you never wrote. Completeness is determined by whoever designs the write path, at design time, and cannot be added later by any amount of signing.

3. Every attestation carries an explicit validity condition.

A control assessed as satisfied on a date is satisfied only until a named change event occurs. Naming those events at attestation time makes evidence expire by design. The alternative is discovering expiry during an audit, which is the same thing arriving at the worst possible moment.

CHAPTER 2

Scope determination

2.0 What counts as an AI system

Before you can classify a system you must know whether it is in scope at all. Article 3(1) defines an AI system as a **machine-based system designed to operate with varying levels of autonomy**, that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, **infers from the input it receives how to generate outputs** such as predictions, content, recommendations or decisions that can influence physical or virtual environments. The Commission published guidelines on this definition in February 2025.

The operative word is **infers**. Two tests follow, and together they resolve most of the arguments teams have about whether something is in scope.

Probably not an AI system	Probably an AI system
A deterministic rules engine where every rule was written by a person and the output is fully traceable to those rules	A model that derives its own decision boundary from data, even where a person reviews the output
A statistical calculation applied to fixed parameters, for example a fee calculator or an actuarial table lookup	A statistical or machine learning model whose parameters were learned from data
A workflow tool that routes items using explicitly coded thresholds	A scoring model that ranks or prioritizes items using learned weights
Basic data processing, search on exact match, or optimization with a fully specified objective and no learned component	A retrieval system whose ranking is learned, or any system built on a foundation model

Where the argument usually happens.

Legacy scoring models. A logistic regression built ten years ago by an actuary, refitted annually on new data, infers from input and is an AI system, whatever it is called internally. The label your organization uses is irrelevant. Systems described as rules engines, decision tables or statistical models must be assessed against Article 3(1) on their mechanics, not their name.

Record the determination. A system excluded from scope because it does not infer needs that reasoning written down as much as an included one does, because the exclusion is the thing an auditor will test.

Four questions decide everything downstream. Answer them per system, not per organization, because a single enterprise routinely holds different roles and different tiers across its portfolio.

1. Does the EU AI Act apply to us at all? 2. What role do we hold for this system? 3. What risk tier does this system fall into? 4. Which other jurisdictions reach this system?

2.1 Question one: does the Act apply

Article 2(1) sets out six categories of actor. The territorial hooks are broader than most organizations assume, and the third one catches companies with no EU presence of any kind.

Art. 2(1)	Who is caught	The practical trigger
(a)	Providers placing AI systems on the Union market or putting them into service, or placing GPAI models on the market	Applies **irrespective of whether the provider is established in the Union or a third country**
(b)	Deployers with place of establishment or located within the Union	An EU subsidiary using a system internally is caught
(c)	Providers and deployers established in a third country **where the output produced by the AI system is used in the Union**	The broadest hook. No EU entity, no EU customer contract, and still in scope if the output lands in the Union
(d)	Importers and distributors of AI systems	
(e)	Product manufacturers placing an AI system on the market together with their product under their own name or trademark	
(f)	Authorized representatives of providers not established in the Union	

Source: Regulation (EU) 2024/1689, Article 2(1). European Commission AI Act Service Desk. Consolidated text at EUR-Lex.

Exclusions worth checking before you conclude you are in scope. Article 2 also carves out scientific research and development where that is the sole purpose, purely personal non-professional use by natural persons, and specified free and open source situations, and it does not affect intermediary liability under Regulation (EU) 2022/2065. Military, defense and national security uses are excluded on their own terms. Verify the paragraph-level wording against EUR-Lex before relying on any exclusion, because each is narrower than its summary.

2.2 Question two: what role do we hold

Roles are per system and they are not interchangeable. The heaviest obligations sit with the provider.

Role	Core position	What it costs you
Provider	Develops the system, or has it developed, and places it on the market or puts it into service under its own name or trademark	Design, risk management, technical documentation, conformity assessment, registration, post-market monitoring
Deployer	Uses the system under its own authority in a professional capacity	Use per instructions, human oversight, input data relevance, monitoring, logging, incident notification
Importer / distributor	Places or makes available a third-country system on the Union market	Verification duties before making available
Product manufacturer	Places an AI system on the market with its product under its own name	Provider obligations for the embedded system

Role migration is the trap. A deployer becomes a provider, and inherits full provider obligations, where it puts its name or trademark on a high risk system, makes a substantial modification to it, or changes its intended purpose such that the system becomes high risk. Fine-tuning a vendor model, rebranding a vendor tool, or deploying a general purpose system into an Annex III use case are all realiztic routes to becoming a provider without a procurement decision having been made. See Article 25.

2.3 Question three: what risk tier

Work the tree in order. Do not skip to Annex III, because prohibition sits above the tier system and applies regardless of tier.

CLASSIFICATION DECISION FLOW

one AI system, one intended purpose. Work in order.

1	**Is it a prohibited practice?** Article 5(1)(a) to (h)	YES	STOP — Withdraw it. This is not a compliance question. See Chapter 3.
		NO	Continue — Go to step 2.
2	**Is it a safety component of, or itself, a product covered by Annex I harmonization legislation requiring third-party conformity assessment?** Article 6(1)	YES	HIGH RISK, ANNEX I — Applies from 2 August 2028.
		NO	Continue — Go to step 3.
3	**Does its intended purpose fall in one of the eight Annex III areas?** Article 6(2). Areas listed below.	YES	Continue — Go to step 4.
		NO	Skip ahead — Go to step 6.
4	**Does the system perform profiling of natural persons?** Absolute bar. No filter is available once this is answered yes.	YES	HIGH RISK — Annex III route. No exemption applies.
		NO	Continue — Go to step 5.
5	**Does it meet one of the four filter conditions, and pose no significant risk of harm to health, safety or fundamental rights?** Article 6(3). Conditions listed below.	YES	NOT HIGH RISK — But Article 6(4) requires the assessment to be documented **before** placing on the market. Undocumented reliance fails.
		NO	HIGH RISK, ANNEX III — Applies from 2 December 2027.
6	**Does it interact with people, generate or manipulate content, or perform emotion recognition or biometric categorization?** Article 50	YES	LIMITED RISK — Transparency duties apply, and are already live.
		NO	MINIMAL RISK — No specific obligations. Baseline governance still recommended.

THE EIGHT ANNEX III AREAS, STEP 3

- Biometrics
- Critical infrastructure
- Education and vocational training
- Employment and worker management
- Essential private and public services
- Law enforcement
- Migration, asylum and border control
- Administration of justice and democratic processes

THE FOUR ARTICLE 6(3) FILTER CONDITIONS, STEP 5

a. Narrow procedural task

b. Improves the result of a previously completed human activity

c. Detects decision-making patterns or deviations, without replacing or influencing the completed human assessment without proper review

d. Performs a preparatory task

Parallel track. If a general purpose AI model is involved, the obligations under Articles 51 to 55 run alongside whatever tier this flow returns. They are not an alternative to it.

Sources: Regulation (EU) 2024/1689, Articles 5, 6, 7, 50 and 51 to 55, and Annexes I and III. The Commission published draft guidelines on the classification of high risk AI systems on 19 May 2026, with consultation closing 23 June 2026. Whether these guidelines are final at the date you read this is on the verify list at Appendix C.

Three classification pitfalls

The filter is read narrowly. A system that makes a value judgment on data relevant to a decision is not performing a narrow procedural task, whatever the interface suggests.

Agentic pipelines are assessed at system level, not component level. Decomposing an agent into individually innocuous steps does not decompose the classification.

Article 6(4) is a pre-market obligation. Reaching a defensible conclusion internally and documenting it afterward fails, because the documentation is the obligation.

2.4 Question four: which other jurisdictions

The EU analysis does not complete scope determination. Record, per system, every market where the system operates or its output lands, then resolve each against the jurisdiction register. The companion workbook carries nineteen regimes that self-flag as in or out of scope. Chapter 5 shows how each resolves back into the same control set, which is the entire point of convergence.

Two cautions for this question specifically. First, voluntary and binding instruments are not equivalent and must not be recorded as though they were. Singapore's Model AI Governance Framework is influential and, for agentic systems, the clearest guidance available, but it is not law and carries no penalty. Chapter 4 states the status of each instrument explicitly. Second, sector regulation stacks on top of horizontal AI law rather than replacing it, and in regulated industries the sector regulator will usually reach you first.

2.5 Scope determination controls

BOARD EXTRACT, LIFTABLE

Scope is settled per system, not per company. For each AI system we record four things: whether the Act reaches it, what role we hold, what tier it falls into, and which other jurisdictions apply. Where we rely on the Article 6(3) exemption to avoid high risk classification, the law requires that reasoning to be documented before deployment. Undocumented reliance is non-compliance regardless of whether the conclusion was correct.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
SCP-001	**Produce a written determination of whether the enterprise falls within Article 2(1), naming which subparagraph is relied on and why**	Scope concluded informally is scope that will be re-litigated under pressure	General Counsel	Signed applicability memo citing the Art. 2(1) subparagraph	EU: Art. 2(1) • NIST: MAP • ISO: 4.1, 4.3	ALL / ALL / EU	Once, then on any change of footprint	New EU entity, new EU customer, new market where output is used, group restructuring
SCP-002	**Record the role held for each inventoried system as provider, deployer, importer, distributor or product manufacturer, with the reasoning**	Provider and deployer obligations differ sharply and the split is per system, not per company	AI Governance Lead	Role register, one row per system, with rationale	EU: Arts. 3, 25, 26 • NIST: MAP • ISO: 4.3	ALL / ALL / EU	On inventory entry, then on change	Rebranding, substantial modification, change of intended purpose, fine-tuning a third-party model
SCP-003	**Screen every inventoried system against Article 5 before any tier analysis, and record the result including for systems you expect to be minimal risk**	Prohibition sits above the tier system, so tier-first screening misses it	General Counsel with AI Governance Lead	Prohibited-practice screening record covering 100 percent of inventory	EU: Art. 5, Recital 44, C(2025) 884 final • NIST: MAP • ISO: 6.1.2	ALL / ALL / EU	Before deployment, then at every inventory sweep	Any change of purpose, user population or deployment context; vendor feature update
SCP-004	**Classify each system by tier and record the rationale, naming the Annex III area or the Annex I legislation relied on**	Regulators challenge classifications, and an undocumented judgment cannot be defended	AI Governance Lead with Legal	Classification record per system with named Annex reference	EU: Arts. 6(1), 6(2), Annexes I and III • NIST: MAP • ISO: 6.1	ALL / ALL / EU	On inventory entry, then on change	Change of intended purpose, new use case, extension to a new population, delegated act amending Annex III
SCP-005	**Where the Article 6(3) filter is relied on, document the assessment before placing on the market or putting into service, and record which of the four conditions applies**	Article 6(4) makes the documentation itself the obligation. Retrospective documentation fails	Provider role owner, countersigned by Legal	Pre-market filter assessment naming the condition and the no-significant-risk conclusion	EU: Arts. 6(3), 6(4) • NIST: MAP • ISO: 6.1.2	P / HR / EU	Before market placement, per system	Any change to intended purpose or autonomy; addition of profiling; system-level recomposition in an agentic pipeline

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
SCP-006	**Confirm and record that no system relying on the Article 6(3) filter performs profiling of natural persons**	Profiling is an absolute bar that no filter condition can overcome	AI Governance Lead with DPO	Profiling determination per filtered system	EU: Art. 6(3); GDPR Art. 4(4) • NIST: MAP • ISO: 6.1.2	P / HR / EU	At filter assessment, then on change	Addition of any personal-data evaluation of performance, behavior, preferences, health or economic situation
SCP-007	**Record every market where each system operates or its output is used, and resolve each against the jurisdiction register**	Scope determined for the EU alone leaves the other regimes unassessed	AI Governance Lead with Legal	Per-system jurisdiction map with in-scope regimes flagged	EU: Art. 2(1)(c) • NIST: MAP • ISO: 4.1, 4.2	ALL / ALL / ALL	On inventory entry, then quarterly	New market entry, new customer geography, change to data residency or output routing
SCP-008	**Determine and record whether a general purpose AI model is involved in each system, and whether your own activity attracts GPAI provider obligations**	Fine-tuning and redistribution can make you a GPAI provider without a decision being taken	Engineering with Legal	GPAI determination per system	EU: Arts. 3, 51 to 55, Annex XIII • NIST: MAP • ISO: A.10	ALL / GP / EU	On inventory entry, then on model change	Model substitution, fine-tuning, redistribution, provider capability change crossing a designation threshold

CHAPTER 3

Red lines: prohibited practices

This is the only chapter where the correct response to a finding is to stop, not to control. Prohibited practices sit above the risk tier system. A system can be minimal risk on every other measure and still be illegal.

Screen the entire inventory against this chapter regardless of tier, and screen before you classify. Prohibitions have applied since **2 February 2025**. Breach attracts the top penalty tier. **Article 99(3)**: up to **35,000,000 euro** or, if the offender is an undertaking, up to **7 percent of total worldwide annual turnover** for the preceding financial year, whichever is higher. Penalties have been enforceable since 2 August 2025 under Article 113.

3.1 The eight prohibitions

Art. 5(1)	Practice	Realiztic enterprise exposure
(a)	Subliminal, purposefully manipulative or deceptive techniques that materially distort behavior and cause or are likely to cause significant harm	**Moderate.** Personalization is not inherently caught. Commission guidance states that personalized advertising is not inherently manipulative. Exposure arises with AI-optimized dark patterns operating below awareness
(b)	Exploiting vulnerabilities of age, disability, or a specific social or economic situation, to distort behavior and cause significant harm	**Moderate.** Relevant to any model that optimizes conversion against financially distressed or otherwise vulnerable segments
(c)	Social scoring: evaluation or classification over time from social behavior or inferred personal characteristics, leading to detrimental treatment that is unrelated to the original context or disproportionate	**Low to moderate.** Legitimate credit and fraud risk modeling is not caught. Document the boundary rather than assuming it
(d)	Predicting the risk of a person committing a criminal offense based solely on profiling or personality traits	**Low** outside law enforcement and some insurance and security contexts
(e)	Untargeted scraping of facial images from the internet or CCTV to build or expand facial recognition databases	**Low**, but check any vendor whose face-matching capability has undisclosed training provenance

Art. 5(1)	Practice	Realiztic enterprise exposure
(f)	**Inferring emotions of a natural person in the areas of workplace and education institutions, except for medical or safety reasons**	**High. This is the one that catches ordinary enterprises.** See 3.2
(g)	Biometric categorization to deduce or infer race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation	**Moderate.** Retail analytics, audience measurement and some HR tooling
(h)	Real-time remote biometric identification in publicly accessible spaces for law enforcement, subject to narrow exceptions	**Low** outside public sector

Sources: Regulation (EU) 2024/1689, Article 5(1)(a) to (h). European Commission AI Act Service Desk. Commission Guidelines on prohibited artificial intelligence practices, C(2025) 884 final, 4 February 2025.

Two scope questions to verify. Whether Article 5(1)(c) social scoring reaches private-sector actors as well as public authorities is summarized inconsistently across secondary sources; check the operative text. And the 2026 Digital Omnibus was reported to extend Article 5 to further practices, including AI-generated child sexual abuse material and non-consensual intimate imagery, with application from 2 December 2026. Both are on the verify list.

3.2 Emotion inference, the prohibition enterprises breach without knowing

Article 5(1)(f) prohibits placing on the market, putting into service for that purpose, or using AI systems **to infer emotions of a natural person in the areas of workplace and education institutions**. Only two exceptions exist: medical reasons and safety reasons.

Workplace is read broadly. Commission guidance treats it as any setting where work is performed, and it covers recruitment and probation, not only established employees. A system assessing the emotional state of an interview candidate is within the prohibition.

The prohibition is narrower than the term emotion recognition system. Article 5(1)(f) does not use that defined term. It prohibits AI systems used to infer emotions. The definition in Article 3 is a useful reference point but is not coextensive with the prohibition.

Who is being assessed decides the outcome. Commission guidance indicates that a call center using voice analysis to track customer emotions is not prohibited by Article 5(1)(f). The same technology pointed at employees, where results reach HR or influence assessment or promotion, is. The technology is not the test. The relationship is.

The medical and safety exceptions are narrow. Pain assessment for non-verbal patients and fatigue detection for safety-critical operation are the archetypes. Wellbeing analytics and engagement scoring are not.

Likely permitted	Likely prohibited
Biometric identity verification for access	Facial expression analysis of employees or students
Behavioral engagement analytics such as click patterns and time on task	Webcam-based attention or engagement inference
Sentiment analysis of customer conversations	Sentiment analysis of employees feeding performance review
Eye tracking for exam integrity that does not infer emotion	Proctoring that infers student affective state

Text-only sentiment analysis is generally not biometric emotion recognition. Where it is applied to employees and influences employment decisions, the boundary is unsettled and specific legal advice is warranted.

3.3 Red line controls

BOARD EXTRACT, LIFTABLE

Eight AI practices are banned outright in the EU and have been since February 2025. Breach carries the maximum penalty of 35 million euro or 7 percent of global turnover. One of the eight catches ordinary companies: inferring the emotions of employees, job candidates or students is prohibited, with only narrow medical and safety exceptions. Facial or voice-based analysis in recruitment, proctoring or workforce analytics is the specific exposure. Doing the same to customers is not prohibited. We screen the entire AI inventory against all eight, not only the systems we consider high risk, because prohibition sits above the risk tier system.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
RED-001	**Screen 100 percent of the inventory against all eight Article 5 practices and record a determination per system, including systems expected to be minimal risk**	Prohibition operates above the tier system, so tier-based screening misses it structurally	General Counsel	Screening register covering every inventoried system, with per-practice determination	EU: Art. 5(1)(a) to (h) • NIST: MAP • ISO: 6.1.2	ALL / ALL / EU	Before deployment, then every inventory sweep	Any new system, change of purpose, new user population, vendor feature release

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
RED-002	**Identify every system that infers emotional or affective state of employees, candidates, or students, and either withdraw it from that use or record the specific medical or safety ground relied on, reviewed by counsel**	The highest-frequency enterprise breach, and it carries the top penalty tier	General Counsel with CHRO	Withdrawal record, or written medical or safety justification with counsel review date	EU: Art. 5(1)(f), Recital 44, C(2025) 884 final • NIST: MAP • ISO: 6.1.2, A.5	ALL / ALL / EU	Immediately, then on every HR or edtech procurement	Any change of purpose or population; extension of results to HR; vendor update adding affect detection
RED-003	**Require every HR, recruitment, proctoring and workforce analytics vendor to state in writing whether their product infers emotional state, and retain the response**	Affect detection arrives as a feature release, not a purchase decision	Procurement with CHRO	Signed vendor attestations, refreshed on renewal	EU: Arts. 5(1)(f), 25 • NIST: MAP • ISO: A.10	D / ALL / EU	At procurement, at renewal, on vendor release notes	Contract renewal, vendor product update, acquisition of the vendor
RED-004	**Document the boundary between your legitimate risk modeling and prohibited social scoring, naming the context limitation and proportionality reasoning**	Credit and fraud modeling is lawful. The defense is the documented boundary, not the intent	Chief Risk Officer with Legal	Boundary memo per scoring system	EU: Art. 5(1)(c) • NIST: MAP • ISO: 6.1.2	ALL / ALL / EU	On deployment, then annually	Extension of a score to a new decision context, addition of inferred characteristics, new data source
RED-005	**Confirm no deployed system performs biometric categorization inferring race, political opinion, trade union membership, religious or philosophical belief, sex life or sexual orientation**	Retail and audience analytics vendors ship this capability by default	CISO with Legal	Per-system biometric categorization determination	EU: Art. 5(1)(g) • NIST: MAP • ISO: 6.1.2	ALL / ALL / EU	On deployment, then every sweep	New analytics deployment, vendor capability change, new camera or sensor estate

CHAPTER 4

The convergence thesis and target operating model

Three instruments, three altitudes, one system. These instruments are not equivalent and must not be recorded as though they were. One is law. One is a certifiable standard. Two are voluntary.

4.1 The four instruments, stated plainly

	EU AI Act	ISO/IEC 42001	NIST AI RMF	Singapore MGF
Status	Binding law. Regulation (EU) 2024/1689	Certifiable international management system standard	Voluntary framework	Voluntary best-practice guidance
Enforced or certified by	National market surveillance authorities; European AI Office for GPAI	Accredited certification bodies, under ISO/IEC 42006	Nobody. No certification exists	Nobody. No certification exists
Applicability trigger	Placing AI on the EU market, being a deployer established in the Union, or output used in the Union (Art. 2(1))	Voluntary adoption. Contractual or customer requirement in practice	Voluntary adoption. Referenced as a safe harbor in Texas TRAIGA	Voluntary adoption. Not tied to Singapore presence
Non-compliance	Up to 35m euro or 7 percent of global turnover (Art. 99)	Loss or refusal of certificate	None	None
What it gives you	The obligation set you must satisfy	Repeatable machinery and third-party assurance	Method for reasoning about risk	Principle-level guidance, strongest available on agentic AI
What it does not	An operating model	Legal compliance, technical method	Certification, legal force	Either of the above

A note on AI Verify. AI Verify is a testing framework and software toolkit from the AI Verify Foundation. It is not a governance framework and is not interchangeable with the MGF. Treat it as a source of test methodology feeding your evaluation evidence, not as a governance instrument. Confirm current scope and version before citing.

A note on Singapore MGF editions. Three exist: the original Model AI Governance Framework for traditional AI, the Model AI Governance Framework for Generative AI, and the Model AI Governance Framework for Agentic AI. The agentic edition launched 22 January 2026 and was updated to version 1.5 on 20 May 2026, refreshed 5 June 2026. Cite the version.

4.2 Three altitudes

NIST reasons. Govern, Map, Measure, Manage. How to think about AI risk, method-rich and technically grounded. It will not certify you and will not make you compliant with anything.

ISO operates. Clauses 4 to 10 and Annex A. How to build machinery that produces governed AI repeatably, and how to prove to a third party that the machinery exists and works. Deliberately non-prescriptive about technical method.

The EU AI Act obliges. What you must achieve and evidence to place AI on the European market lawfully. It specifies outcomes, not operating models.

Singapore enriches the reasoning layer and, for agentic systems, is currently the most specific guidance available anywhere.

4.3 Why one control set rather than four programs

The convergence test is whether a single artifact can discharge obligations under multiple instruments simultaneously. It can, and this is the whole basis of the document.

One artifact	Satisfies
Immutable inference log emitted by the enforcement point, recording principal, model version, input reference, output reference, timestamp	EU Arts. 12 and 19 record-keeping; NIST MEASURE; ISO Clause 8 and A.6; sector audit trail expectations
Pre-deployment risk and impact assessment with named residual risk acceptance	EU Arts. 9 and 27; NIST MAP and MANAGE; ISO Clauses 6.1 and 6.1.3; SG risk assessment
Per-agent identity bound to an authorizing human, with scoped tool permissions	EU Arts. 14 and 15; NIST GOVERN and MANAGE; ISO A.9; SG MGF Agentic dimensions 2 and 3

Built once, mapped many times. Chapter 5 sets out the full crosswalk. Chapters 6 to 14 carry the controls that produce these artifacts.

4.4 Target operating model

Component	What it is	Accountable	Failure signature
Single accountable owner	One named individual owning the program end to end, with budget and the standing to stop a deployment	Named AI Governance Lead or CAIO	Governance distributed across a committee, so no decision is anyone's to make
One control set	The crosswalk in Chapter 5, adopted as the enterprise control framework rather than one framework per instrument	AI Governance Lead	Parallel control libraries per instrument, each partially maintained
Enforcement points that emit	Technical chokepoints, principally the inference gateway and the deployment gate, that produce evidence as a byproduct of enforcing	Head of Engineering with CISO	Evidence produced by a reporting layer that systems can bypass
An evidence register with expiry	One index mapping artifact to control to obligation, where every entry carries a validity condition	AI Governance Lead	Evidence that is complete on the day of the audit and silently stale thereafter
A review rhythm	Quarterly reassessment of changed systems, monthly inventory sweep, annual full re-run	AI Governance Lead	An assessment completed once and filed

The gate must be the only path.

A deployment gate that a pipeline can route around is a suggestion. A gate that emits the approval record as a condition of releasing is a control.

Ownership of the write path is ownership of completeness.

Whoever designs where records are written determines what can never be evidenced later. That decision is made once, early, usually by an engineer who was not told it was a governance decision. Chapter 9 makes it one.

CHAPTER 5

The unified control set

Each row is a governance capability the enterprise needs. The instrument columns show where that capability is addressed. The evidence column names the artifact that discharges all of them at once.

5.1 How to read the crosswalk

Build to the evidence column, not to the instrument columns. The instrument columns tell you why a capability exists. The evidence column tells you what to produce. Teams that build to the instrument columns build the same control several times.

A capability is not a control. Each row is a heading under which one or more executable controls sit. The controls are in Chapters 6 to 14. This chapter is the map, those chapters are the territory.

Voluntary columns do not create obligations. NIST and Singapore references indicate where those frameworks address the capability. Only the EU column carries legal force, and only where the Chapter 2 scope analysis puts you in its reach.

EU references are Regulation (EU) 2024/1689 articles. ISO references are ISO/IEC 42001 clauses and Annex A groups, by number and intent only. NIST references are AI RMF functions. Singapore references are Model AI Governance Framework dimensions and are indicative mappings rather than IMDA statements.

5.2 The crosswalk

#	Capability	EU AI Act	NIST	ISO/IEC 42001	Singapore MGF	Evidence that satisfies all four	Owner	P
1	**AI system inventory**	Arts. 16, 49, 71	MAP	4.3, A.4	Accountability	Living inventory with per-system tier, role, owner, status	AI Gov Lead	P1
2	**Risk assessment**	Art. 9	MAP, MEASURE	6.1, A.5	Risk assessment	Documented assessment per system, refreshed on trigger	CRO	P1
3	**AI impact assessment**	Art. 27	MAP	6.1, A.5	Accountability	Impact assessment covering rights, safety, affected groups	Risk with Legal	P1
4	**Risk classification**	Art. 6, Annex III	MAP	6.1	Risk assessment	Tier determination with written rationale per system	Legal with Risk	P1
5	**Data governance**	Art. 10	MAP, MEASURE	A.7	Data	Lineage, quality, provenance and minimization records	CDO	P1

#	Capability	EU AI Act	NIST	ISO/IEC 42001	Singapore MGF	Evidence that satisfies all four	Owner	P
6	**Human oversight**	Art. 14	GOVERN, MANAGE	5.3, A.9	Agentic D2	Oversight design per system plus evidence gates are exercised	Business owner	P1
7	**Transparency to users**	Arts. 13, 50	MAP, MANAGE	A.8	Content provenance	Disclosures, interaction notices, content labeling	Product	P2
8	**Technical documentation**	Art. 11, Annex IV	MAP, MEASURE	7.5	Trusted development	Current technical file per high risk system	Engineering	P2
9	**Logging and traceability**	Arts. 12, 19	MEASURE, MANAGE	8.1, A.6	Agentic D3	Immutable logs: principal, model version, input, output, time	Eng with Security	P1
10	**Monitoring and post-market**	Art. 72	MEASURE, MANAGE	9.1	Incident reporting	Continuous monitoring against defined thresholds	Eng with Ops	P2
11	**Model validation and testing**	Art. 15	MEASURE	8.1, A.6	Testing and assurance	Validation results, coverage, acceptance criteria	Data Science	P2
12	**Incident management**	Art. 73	MANAGE	10.2, A.6	Incident reporting	AI-specific incident process, records, notifications	Security with Risk	P1
13	**Security, AI-specific**	Art. 15	MEASURE, MANAGE	A.6	Security	Adversarial test results, guardrail validation	CISO	P1
14	**Bias and fairness**	Arts. 10, 15	MEASURE	A.5, A.6	Data	Bias testing across relevant groups, mitigation records	DS with Risk	P2
15	**Privacy**	Art. 10, and GDPR	MAP, MEASURE	A.7, and 27701	Data	DPIA where applicable, data subject handling	DPO	P1
16	**Explainability**	Arts. 13, 14, 86	MEASURE	A.8	Trusted development	Explainability approach matched to audience	Data Science	P3
17	**Third party and vendor AI**	Art. 25	GOVERN, MAP	A.10	Accountability	Vendor assessments, contractual controls, provenance	Procurement with Risk	P2
18	**Supplier management**	Art. 25	GOVERN	A.10	Accountability	Due diligence and ongoing monitoring records	Procurement	P2
19	**Model and system retirement**	Lifecycle	MANAGE	A.6	Trusted development	Decommissioning process, data handling, records	Engineering	P3
20	**Drift detection**	Arts. 15, 72	MEASURE	9.1, A.6	Incident reporting	Drift metrics, thresholds, retraining triggers	DS with Ops	P2
21	**Red teaming**	Art. 15	MEASURE	A.6	Testing and assurance	Adversarial campaigns, findings, remediation	Security	P2
22	**Testing and QA**	Arts. 15, 17	MEASURE	8.1	Testing and assurance	Test strategy, coverage, results, gates	Eng with QA	P2
23	**Board reporting**	Art. 26	GOVERN	5.1, 9.3	Accountability	Regular board-level AI risk and compliance reporting	AI Gov Lead	P2
24	**Training and awareness**	Art. 4	GOVERN	7.2, 7.3	Accountability	Role-based training records, literacy program	HR with AI Gov	P1
25	**Competence**	Art. 4	GOVERN	7.2	Accountability	Competence definitions per role, evidence held	HR	P2
26	**Business continuity**	Art. 15	MANAGE	A.6, and 22301	Incident reporting	Continuity plans covering AI system failure	Ops with Risk	P3

#	Capability	EU AI Act	NIST	ISO/IEC 42001	Singapore MGF	Evidence that satisfies all four	Owner	P
27	**Policy framework**	Art. 17	GOVERN	5.2, A.2	Accountability	Approved AI policy set, enforced and reviewed	AI Gov Lead	P1
28	**Roles and accountability**	Arts. 22, 26	GOVERN	5.3, A.3	Accountability	RACI with single accountability per activity	AI Gov Lead	P1
29	**Conformity assessment**	Arts. 43, 47, 48	Not addressed	9.2, and 42006	Not addressed	Conformity records, CE marking evidence	Compliance	P2
30	**Continuous improvement**	Art. 72	MANAGE	10.3	Incident reporting	Nonconformity handling, corrective actions	AI Gov Lead	P3
31	**GPAI and foundation models**	Arts. 51 to 55	MAP, MEASURE	A.10, A.6	GenAI	Model documentation, usage controls, provider evidence	Eng with Legal	P2
32	**Agentic AI governance**	Arts. 14, 15	GOVERN, MANAGE	A.9, A.6	Agentic D1 to D4	Agent identity, least privilege, kill switch, oversight	CISO with Eng	P2

5.3 Three notes on the gaps

Two rows show not addressed in the voluntary columns. This is informative rather than a defect. Conformity assessment is a legal construct with no equivalent in a voluntary framework, and NIST and Singapore correctly do not attempt one. Where a capability appears in only one instrument column, that instrument is your sole authority for it.

Capability 32 is the row where the voluntary instrument is ahead of the law. The EU AI Act reaches agents through Articles 14 and 15, written before agentic deployment was common. Singapore's agentic framework is more specific. Where they differ in granularity, follow Singapore for design and the EU AI Act for what you must evidence.

Capability 9 is the row that determines whether most of the others can be evidenced at all. Chapter 9 treats it as an architecture problem rather than a logging problem.

CHAPTER 6

AI system inventory and registration

ENFORCEMENT POINT FOR THIS DOMAIN

The deployment gate. An inventory maintained by survey drifts within weeks. An inventory that a system is entered into as a precondition of being deployed cannot drift, because the gate is the only path to production.

6.1 Why this chapter comes before every other control chapter

Every downstream obligation is per system. Risk assessment, classification, oversight design, logging, monitoring and evidence all require a defined subject. An enterprise that cannot enumerate its AI systems cannot perform any of them completely, and completeness is the failure mode auditors probe first.

The gap between the official inventory and the real one has three sources: AI arriving as a feature in software already licensed, AI procured below approval thresholds by individual teams, and AI reached through personal accounts. The first is the largest and the least visible.

6.2 What a compliant inventory record contains

Field	Why it is required
System identifier and name	Stable reference for every downstream artifact
Intended purpose, as documented	Determines classification. Article 6 turns on intended purpose, not capability
Business owner and accountable individual	Named person, not a team
Our role: provider, deployer, importer, distributor, product manufacturer	Determines which obligations attach. Per system, per Chapter 2

Field	Why it is required
Risk tier and the rationale	Drives everything. Rationale is required, not optional
Annex reference relied on	Names the specific Annex III area or Annex I legislation
Underlying models and versions, including third party	Required for provenance, and for revalidation when a vendor swaps a model
Data sources and categories	Links to data governance and privacy obligations
Lifecycle stage and deployment status	Distinguishes production from pilot, which changes obligations
Jurisdictions where operated or output used	Drives the multi-jurisdiction analysis
Autonomy level and action space, if agentic	Determines the agent control set in Chapter 13
Last assessment date and next review trigger	Makes staleness visible
Evidence location	Links the register to the evidence index in Chapter 9

6.3 Registration

Providers of high risk systems listed in Annex III are required to register in the EU database before placing on the market or putting into service, with certain deployers also subject to registration duties. Articles 49 and 71 govern this. The operational mechanics of the database, including precisely who registers what and when, have been fluid and are on the verify list.

Do not treat registration as a filing task. The information required is a subset of what the technical file already contains, so an organization whose documentation is current can register quickly, and one whose documentation is retrospective cannot.

6.4 Inventory and registration controls

BOARD EXTRACT, LIFTABLE

We cannot govern what we have not enumerated, and every obligation in this program is per system. Our inventory is enforced at the deployment gate rather than collected by survey, so it cannot silently decay. The largest source of unregistered AI is not shadow procurement but AI features added by vendors to software we already license, which is why we screen the software register quarterly rather than relying on purchase records.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
INV-001	**Run a discovery sweep across procurement records, expense claims, SaaS admin consoles, network egress and identity logs, and record the method and coverage achieved**	A sweep that samples cannot evidence completeness, and completeness is the point	AI Gov Lead with CISO	Sweep report naming sources searched, coverage, and systems found	EU: Arts. 16, 26 • NIST: MAP • ISO: 4.3	ALL / ALL / ALL	Initial, then monthly	New SaaS agreement, new business unit, acquisition, expense category change
INV-002	**Require every AI system to be entered in the register as a precondition of passing the deployment gate, and configure the gate to refuse release where no register entry exists**	Inventory maintained by survey decays. Inventory enforced by the gate cannot	Head of Engineering	Gate configuration under version control, plus a bypass test showing no release path skips the check	EU: Art. 17 • NIST: GOVERN • ISO: 8.1	ALL / ALL / ALL	At gate implementation, then re-tested on pipeline change	New deployment pipeline, new environment, change to release tooling, exception granted to any team
INV-003	**Record the documented intended purpose for each system in the words used to specify it, and record separately any purpose the system is capable of but not authorized for**	Classification turns on intended purpose. A purpose written after classification is circular	Business owner with AI Gov Lead	Intended purpose statement per system, dated and version controlled	EU: Arts. 3, 6 • NIST: MAP • ISO: 4.3, 6.1	ALL / ALL / EU	On entry, then on change	Any expansion of use, new user group, new integration, feature release changing capability
INV-004	**Name one accountable individual per system and record the date they accepted accountability in writing**	A team named as owner means no person accepted anything	AI Governance Lead	Signed accountability record per system, with name and date	EU: Arts. 22, 26 • NIST: GOVERN • ISO: 5.3	ALL / ALL / ALL	On entry, then on change	Departure or role change of the named individual, transfer of the system between units
INV-005	**Screen every SaaS product in the software register for AI features added since procurement, and record the result per product**	Embedded AI arrives by release note, not by purchase order. This is the largest source of unregistered AI	Procurement with AI Gov Lead	Screening record per SaaS product with date and finding	EU: Arts. 25, 26 • NIST: MAP • ISO: A.10	D / ALL / ALL	Quarterly, and on vendor release notes	Vendor release adding AI capability, contract renewal, vendor acquisition

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
INV-006	**For each system, record every jurisdiction in which it operates or its output is used, and flag the regimes that apply**	Article 2(1)(c) reaches systems with no EU entity behind them. Unrecorded geography is unassessed obligation	AI Gov Lead with Legal	Per-system jurisdiction map with in-scope regimes	EU: Art. 2(1) • NIST: MAP • ISO: 4.1, 4.2	ALL / ALL / ALL	On entry, then quarterly	New market, new customer geography, change to output routing or data residency
INV-007	**Register high risk Annex III systems in the EU database before placing on the market or putting into service, and retain the submission record**	Registration is a precondition, not a follow-up task	Compliance	Registration submission and confirmation records	EU: Arts. 49, 71 • ISO: 7.5	P / HR / EU	Before market placement	Change to registered particulars, change of intended purpose, substantial modification
INV-008	**Record autonomy level and permitted action space for every agentic system, using a defined scale**	Agent risk is set by what it can do without a human, and that is not captured by a risk tier	CISO with Engineering	Agent register with autonomy level and action space per agent	EU: Arts. 14, 15 • NIST: MAP • ISO: A.9 • SG: Agentic D1	ALL / ALL / ALL	On entry, then on any permission change	New tool granted, scope expansion, autonomy increase, addition to a multi-agent pipeline

CHAPTER 7

Risk management and classification across the lifecycle

ENFORCEMENT POINT FOR THIS DOMAIN

The deployment gate and the change pipeline. A risk assessment produced on request is a document. A risk assessment that the gate requires before release, and that the change pipeline invalidates on a triggering event, is a control.

7.1 What Article 9 actually requires

Article 9(1): a risk management system shall be **established, implemented, documented and maintained** in relation to high-risk AI systems. Four verbs, four distinct failures. Most programs establish and document. Fewer implement. Very few maintain.

Article 9(2): the risk management system shall be understood as **a continuous iterative process planned and run throughout the entire lifecycle** of a high-risk AI system, **requiring regular systematic review and updating**.

A point-in-time assessment does not satisfy Article 9, however thorough. The obligation is to run a process, not to hold a document. An organization with excellent assessments and no review cadence is non-compliant in a way that no amount of assessment quality repairs.

The risk management system is the layer that forces the other obligations into one process. Data governance under Article 10, transparency under Article 13, oversight under Article 14 and robustness under Article 15 are each measures that Article 9 coordinates. Findings from Article 10 bias examination are not engineering notes. They are inputs to the Article 9 process.

The penalty tier, now verified. Article 9 sits below the Article 5 tier. Under **Article 99(4)**, non-compliance with the obligations of operators attracts fines up to **15,000,000 euro or 3 percent** of total worldwide annual turnover, whichever is higher. The paragraph enumerates the obligations of providers under Article 16, authorized representatives under Article 22, importers under Article 23, distributors under Article 24, deployers under Article 26, notified bodies, and the transparency obligations under Article 50. The Article 9 to 15 requirements are enforced through Article 16, which is how a risk management failure reaches this tier.

Article 99(5) sets a third tier of up to **7,500,000 euro or 1 percent** for supplying incorrect, incomplete or misleading information to notified bodies or authorities in reply to a request. Do not repeat the 35 million euro figure in relation to Article 9. That tier belongs to Article 5 alone.

The SME reversal, Article 99(6).

For SMEs and start-ups, each fine is capped at the LOWER of the fixed amount and the percentage, which is the exact reverse of the whichever is higher rule that applies to everyone else. It is written into the Regulation and rarely quoted, and it materially changes exposure modeling for smaller undertakings. GPAI model provider fines sit separately under Article 101, enforced by the Commission and the AI Office.

7.2 The classification to obligation chain

Tier determined in Ch 2	What switches on
Prohibited	Nothing. The system is withdrawn. See Chapter 3
High risk, Annex I or III	Arts. 9, 10, 11, 12, 13, 14, 15, 17, 43, 49, 72, 73. The full regime
Limited risk	Art. 50 transparency only
Minimal risk	No specific obligations. Baseline governance recommended
GPAI model involved	Arts. 51 to 55 run in parallel with whatever tier applies

Classification must be settled before risk assessment scope can be set, and reclassification must reopen the assessment. Teams that assess first and classify second produce assessments scoped to the wrong obligation set.

7.3 Reassessment triggers

Trigger type	Events
Model	Retraining, fine-tuning, version upgrade, vendor-initiated model substitution, change of provider
Purpose and scope	New use case, new user population, new geography, extension to a new decision type, increase in autonomy
Performance	Accuracy or fairness metric breaching threshold, drift detection alert, sustained degradation
Event	Any incident on the system, any incident on a comparable system, complaint pattern, near miss
External	Regulatory change, delegated act amending Annex III, new guidance, harmonized standard publication
Periodic	Defined cycle regardless of the above, at minimum annually for high risk

Vendor-initiated model substitution deserves separate emphasis. It is the trigger most likely to fire without anyone in the organization taking a decision, and the one most likely to be missed because it arrives as a release note.

7.4 Residual risk acceptance

Article 9(5) requires that residual risk be judged acceptable. A judgment implies a judge. Three failures recur: the acceptance is unnamed, so no individual is on record; undated, so it cannot be tied to the system state it applied to; and unconditional, so it survives changes that should have voided it.

7.5 Risk management controls

BOARD EXTRACT, LIFTABLE

The law requires a risk management process that runs continuously across the life of each high risk system, not an assessment performed once before launch. Our assessments are enforced at the deployment gate, so a system cannot reach production without one. Reassessment is triggered by defined events, including vendor model substitutions that arrive without a decision on our part. Every residual risk acceptance carries a named individual, a date, and the conditions that void it.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
RSK-001	**Publish a written risk assessment methodology stating the scoring scales, the evaluation criteria, and who may accept risk at each level**	Assessments produced without a common method cannot be compared, aggregated or defended	Chief Risk Officer	Approved methodology document, version controlled	EU: Art. 9(1), 9(2) • NIST: MAP, MEASURE • ISO: 6.1 • ISO/IEC 23894	ALL / HR / EU	Once, then annually	Change to risk appetite, new system category, regulatory change altering assessment scope
RSK-002	**Complete a documented risk assessment for each high risk system before it passes the deployment gate, and configure the gate to refuse release without one**	Post-deployment assessment records a decision already taken	CRO with Head of Engineering	Per-system assessment, plus gate configuration showing the check is mandatory	EU: Art. 9(1) to 9(4) • NIST: MAP, MEASURE • ISO: 6.1, 8.1	P, D / HR / EU	Before every release	New release, model change, purpose change, pipeline change bypassing the gate
RSK-003	**Define the full trigger list that reopens a risk assessment, publish it, and instrument the model registry and change pipeline to fire on each machine-detectable trigger**	A trigger list that relies on someone remembering is not a control	CRO with Engineering	Published trigger list plus instrumentation evidence showing automated firing	EU: Art. 9(2) • NIST: MANAGE • ISO: 6.1, 9.1	ALL / HR / EU	At implementation, then verified quarterly	New trigger type identified, registry or pipeline change, new deployment path
RSK-004	**Record residual risk acceptance for each high risk system with the accepting individual's name, the date, the risk level accepted, and the conditions that void the acceptance**	An unnamed, undated, unconditional acceptance is not an acceptance	CRO	Signed residual risk acceptance record per system	EU: Art. 9(5) • NIST: MANAGE • ISO: 6.1.3	ALL / HR / EU	At deployment, then on each trigger	Departure of the accepting individual, any listed trigger firing, risk level exceeding the accepted band
RSK-005	**Route every Article 10 bias finding, every drift alert and every incident into the risk register as an entry, rather than holding them in engineering or security systems**	Article 9 is the process that coordinates the other obligations. Findings held elsewhere never reach it	CRO with Engineering	Register entries traceable to their originating finding, with source system reference	EU: Arts. 9(2)(c), 10(2)(f) • NIST: MEASURE, MANAGE • ISO: 6.1, 10.2	ALL / HR / EU	Continuous	New finding source added, change to monitoring or incident tooling

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
RSK-006	**Assess and record, for each high risk system, whether it is likely to have an adverse impact on persons under 18 or other vulnerable groups**	Article 9 requires this specific consideration and it is routinely omitted	CRO with Legal	Vulnerable group impact determination per system	EU: Art. 9 • NIST: MAP • ISO: 6.1, A.5	ALL / HR / EU	At assessment, then on population change	Extension to a new user population, change of distribution channel, age gating change
RSK-007	**Test each high risk system against previously defined metrics and thresholds appropriate to its intended purpose, before market placement, and retain the results**	Testing against thresholds chosen after seeing the results is not testing	Engineering with QA	Test results against pre-agreed thresholds, dated before release	EU: Arts. 9, 15 • NIST: MEASURE • ISO: 8.1	P / HR / EU	Before market placement, then on retraining	Retraining, model substitution, threshold revision, change of intended purpose
RSK-008	**Reassess every high risk system on a defined periodic cycle regardless of whether a trigger has fired, and record the review even where nothing changed**	Regular systematic review is required by Article 9(2) independently of event-driven review	CRO	Periodic review records including no-change reviews	EU: Art. 9(2) • NIST: MANAGE • ISO: 9.1, 9.3	ALL / HR / EU	At minimum annually	Cycle elapsing, change to review cadence policy

CHAPTER 8

Data governance, provenance and quality

ENFORCEMENT POINT FOR THIS DOMAIN

The data pipeline and the retrieval layer. Provenance asserted in a document is a claim. Provenance recorded by the pipeline that ingested the data is evidence. Permission enforced at retrieval is a control; permission checked at the application layer is a preference.

8.1 What Article 10 requires

Article 10(1) applies to high risk systems that use techniques involving the **training of AI models with data**, and requires that they be developed on the basis of **training, validation and testing data sets** meeting the quality criteria in paragraphs 2 to 5.

Article 10(2) requires those data sets to be subject to data governance practices appropriate to the intended purpose, concerning in particular design choices, data collection processes and the origin of data, and in the case of personal data the **original purpose of collection**, data preparation operations such as annotation, labeling, cleaning, updating, enrichment and aggregation, the assumptions made about what the data represents, assessment of availability, quantity and suitability, **examination for possible biases** likely to affect health and safety, negatively affect fundamental rights or lead to prohibited discrimination, measures to detect, prevent and mitigate those biases, and identification of data gaps or shortcomings that prevent compliance.

Article 10(3) requires data sets to be **relevant, sufficiently representative, and to the best extent possible free of errors and complete** in view of the intended purpose, with appropriate statistical properties including as regards the persons or groups the system is intended to be used on.

Article 10(4) requires account to be taken of characteristics particular to the geographical, contextual, behavioral or functional setting. Article 10(5) permits providers exceptionally to process special categories of personal data strictly for bias detection and correction, subject to safeguards. Article 10(6) applies paragraphs 2 to 5 to testing data sets only, where the system does not involve training.

Two observations. First, to the best extent possible is a qualified standard, but it requires you to show what was possible and what you did. Second, Article 10 is largely an ex ante obligation while Article 9 is continuous, which leaves a gap: continuous risk assessment without corresponding continuous data governance. Close it deliberately.

8.2 The three data governance failures that matter

Lineage that cannot be reconstructed.

Without dataset versioning tied to model versions, you cannot state which data produced which model. That makes reproduction impossible, bias source tracing impossible, and defense of a contested decision impossible.

The retrieval permission gap.

In retrieval-augmented systems, if the retrieval layer does not enforce the permissions of the source documents, the system will answer using material the user was never entitled to see. This is the most common serious data failure in production generative systems, and it is not detected by any control that examines training data.

Deletion that leaves derivatives intact.

Deleting a source record while its embeddings, cached retrievals, fine-tuning extracts and logs persist is not deletion. Retention schedules written for source data and silent on derivatives fail on first contact with a deletion request.

8.3 Data governance controls

BOARD EXTRACT, LIFTABLE

The law sets quality obligations on the data behind high risk AI: it must be relevant, representative, and as free of errors and as complete as the purpose demands, with its origin documented and its biases examined. Two failures carry the most exposure. First, if we cannot state which data produced which model, we cannot defend any decision that model made. Second, in retrieval-based systems, permissions must be enforced where documents are fetched rather than in the application, or the system will answer using material the user was never entitled to see. Both are architectural choices, not policy choices.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
DAT-001	**Record, for each training, validation and testing data set, its origin, collection process, and where personal data is involved the original purpose of collection**	Article 10(2)(b) requires origin and original collection purpose specifically, and it cannot be reconstructed later	Chief Data Officer	Provenance record per data set, emitted by the ingestion pipeline	EU: Art. 10(2)(a), 10(2)(b) • NIST: MAP • ISO: A.7	P / HR / EU	At ingestion	New data source, change to collection method, re-purposing of an existing set

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
DAT-002	**Version every data set and bind the version to the model version trained on it, so any model in production resolves to an exact data state**	Without this binding, no claim about training data can be substantiated	Engineering with CDO	Dataset version register with model-to-dataset bindings	EU: Arts. 10(2), 11 • NIST: MAP, MEASURE • ISO: A.7, 7.5	P / HR / ALL	At every training run	Retraining, data refresh, model substitution, pipeline change
DAT-003	**Record all data preparation operations applied, including annotation, labeling, cleaning, updating, enrichment and aggregation**	Article 10(2)(c) names these operations explicitly. Preparation is where bias is most often introduced and least often recorded	CDO	Preparation log per data set, generated by the pipeline	EU: Art. 10(2)(c) • NIST: MAP • ISO: A.7	P / HR / EU	At every preparation run	New preparation step, change to annotation guidance, change of labeling vendor
DAT-004	**Examine each data set for biases likely to affect health and safety, adversely affect fundamental rights, or lead to prohibited discrimination, and record both the examination and the mitigations applied**	Article 10(2)(f) and (g) require examination and mitigation as separate obligations	CDO with Data Science	Bias examination report and mitigation record per data set	EU: Arts. 10(2)(f), 10(2)(g) • NIST: MEASURE • ISO: A.5, A.7	P / HR / EU	Before training, then on data refresh	Data refresh, population change, deployment in a new geography or setting
DAT-005	**Record identified data gaps and shortcomings that prevent compliance, with the plan to address them**	Article 10(2)(h) requires the gaps to be identified, not only the strengths	CDO	Data gap register with remediation plan	EU: Art. 10(2)(h) • NIST: MAP • ISO: A.7	P / HR / EU	At assessment, then on refresh	New gap identified, remediation completed, change of intended purpose
DAT-006	**Enforce source document permissions at the retrieval layer, and test by attempting retrieval as a low-privilege principal against a document that principal cannot open**	Application-layer permission checks do not prevent the model answering from restricted material	Head of Engineering with CISO	Retrieval permission test report showing the low-privilege attempt returned nothing	EU: Art. 10, and GDPR • NIST: MEASURE • ISO: A.7	ALL / ALL / ALL	At deployment, then quarterly	New retrieval source, index rebuild, permission model change, new tenant
DAT-007	**Extend retention and deletion schedules to cover derived artifacts including embeddings, vector indexes, cached retrievals, fine-tuning extracts and prompt logs**	Deleting the source and retaining the derivatives is not deletion	CDO with DPO	Retention schedule naming each derivative class, plus a deletion test record	EU: Art. 10, and GDPR Art. 5 • NIST: MANAGE • ISO: A.7	ALL / ALL / ALL	At implementation, then annually	New derivative artifact type, new vector store, change of retention period
DAT-008	**Where special category personal data is processed for bias detection and correction, record the necessity justification and the safeguards applied**	Article 10(5) permits this exceptionally and conditionally. The record is what makes it lawful	DPO with CDO	Necessity and safeguards record per processing activity	EU: Art. 10(5) • NIST: MAP • ISO: A.7, and 27701	P / HR / EU	Before processing	Change to the bias testing method, new special category introduced, change of legal basis

CHAPTER 9

Evidence architecture

This chapter is about enforcement points, so it has no separate one. Its subject is the question every other chapter depends on: when you say a control operated, what makes that statement true.

9.1 Three failures usually treated as one

Failure	The question it answers badly	What it looks like
Fidelity	Does the record describe what actually happened?	The record reflects what a reporting layer was told, not what the system did
Completeness	Are these all the records there should be?	The records held are internally consistent and silently partial
Currency	Is this still true?	The record was accurate on the date it was made and describes a system that no longer exists

These require different remedies. Fidelity is fixed by moving emission to the enforcement point. Completeness is fixed at the write path and cannot be fixed anywhere else. Currency is fixed by attaching validity conditions at the time of attestation.

Most evidence programs invest in integrity, which addresses none of the three. Integrity proves that the records you hold were not altered after the fact. It is necessary and it is not sufficient.

9.2 Principle one: evidence is emitted by the enforcement point

Definition. An enforcement point is a component that must be traversed for the action to occur. An observer is a component that watches actions occur. Both can produce records. Only one produces evidence.

The test. Can the action occur without traversing this component? If yes, it is an observer. And if the emitter fails, does the action fail? If the action proceeds when recording fails, the record is missing precisely when the system is under stress. Fail-open logging guarantees that your evidence is thinnest during exactly the incidents you will later be asked about.

Action to evidence	Enforcement point	Common observer used instead
A model was invoked, by whom, with what	Inference gateway or model proxy	Application logging in the calling service
A system reached production	Deployment gate in the release pipeline	Change ticket in the service desk
An agent took an action	Tool broker or function-call authorizer	Agent framework trace, written by the agent
Context was retrieved	Retrieval layer	Application query log
A human approved a decision	Approval service that issues the token the action requires	Screenshot, email, or a field in a case record
Access was granted	Identity provider at token issuance	Periodic access review export

The agent case deserves emphasis. An agent framework that writes its own trace is the purest form of self-reported evidence. The agent records what it believes it did. If the agent is compromised, misconfigured, or simply has a defect in its trace logic, the record follows the defect. The enforcement point for an agent action is the component that authorizes the tool call, because the action cannot occur without it.

9.3 Principle two: completeness is not tamper evidence

These properties are orthogonal. Tamper evidence answers: were the records I hold altered? It is solved by hash chaining, signing, append-only storage and transparency logs. These are mature and they work.

Completeness answers: are these all the records there should be? **No cryptographic control addresses this.** A perfectly signed, immutably stored, hash-chained log that is missing forty percent of events is cryptographically sound and evidentially worthless. The signature attests to what was written. It is silent on what was not.

This matters because integrity is the property that gets procured. It is visible, it is a product category, and it demonstrates well. Completeness is invisible, has no product category, and is decided by an engineer choosing where to put a function call.

Who decides completeness. Whoever designs the write path. Specifically, whoever decides which code paths call the emitter and whether any reach the resource without doing so; whether emission is synchronous or queued; whether emission failure blocks or is swallowed; whether records carry a sequence so that absence is detectable; and what happens when the buffer fills, the disk fills, or the destination is unreachable.

Each of these is made once, early, usually during implementation, usually by someone who was not told it was a governance decision. By the time governance sees the system, the completeness ceiling is already set, and no amount of downstream control raises it.

Four controls that actually establish completeness

Sequence numbering with gap detection. Records carry a monotonic sequence per emitter. A gap is then detectable rather than invisible. Without this, missing records are indistinguishable from periods of inactivity.

Heartbeat records. An emitter that has produced nothing is ambiguous between nothing happened and recording stopped. A periodic null record disambiguates it. This is cheap and almost never done.

Independent reconciliation. Compare your record count against a count produced by a party with no interest in your evidence. Your model provider bills per call or per token, so the invoice is an independent count of inferences. If your log holds fewer inferences than you were billed for, you have a completeness gap and you can quantify it to the record. This is the single strongest completeness control available and it costs almost nothing.

Bypass testing. Enumerate every path to the resource and attempt the action through each, confirming emission. Repeat when the topology changes. An architecture diagram is a claim about paths; a bypass test is evidence about paths.

9.4 Principle three: attestations carry validity conditions

An attestation is a statement about a system state at a moment. Systems change continuously. The attestation does not know this and will assert itself indefinitely. Under current practice, expiry is discovered rather than designed. Someone examines the evidence during an audit, notices the model version has moved three times since the test, and the finding is raised. The evidence did not become false at that moment. It became false months earlier, and the organization operated on it in the interval.

Element	What it states	Example
Scope	Exactly what the attestation covers	Inference gateway v4.2 in production, EU region
Basis	What was examined to reach it	Bypass test across six enumerated paths, 12 Aug 2026
Voiding events	Enumerated changes that make it false	New endpoint, gateway version change, new direct-calling client, log destination change
Watcher	What detects each voiding event	Service registry alert, CI pipeline hook, IAM change stream
Hard expiry	Maximum validity irrespective of events	12 months
Revalidation	What must be redone, and whether partially	Re-run bypass test on affected paths only

Class	Typical voiding events
Configuration	Threshold changed, guardrail setting altered, retention period reduced
Version	Model, prompt, gateway, dependency, agent framework

Class	Typical voiding events
Scope	New user population, new geography, new decision type, autonomy increase
Topology	New endpoint, new integration, new client, new retrieval source, new tenant
Personnel	Departure or role change of the accountable individual
External	Regulatory change, new guidance, delegated act, harmonized standard published
Time	Hard expiry regardless of the above

The watcher requirement. A voiding event that nothing detects will not fire. An attestation naming voided by model version change where no process watches the model registry is a documented intention, not a control. Prefer conditions where the detection already exists: version registries, configuration management, identity systems, change pipelines.

A voiding event that fires does not mean the control failed. It means the evidence has expired and revalidation is due. Treating expiry as a finding creates an incentive to write weak validity conditions, which defeats the mechanism. Treat expiry as scheduled work.

9.5 The evidence register

Field	Purpose
Artifact identifier and location	Retrieval within the target window
Controls satisfied	Many-to-many. One artifact typically discharges several
Obligations satisfied	The four instrument references, so one artifact maps to all applicable regimes
Emitting component	Which enforcement point produced it. Blank means it was narrated, which is itself a finding
Completeness basis	Sequence, heartbeat, reconciliation, or none
Attested on, by	Date and named individual
Validity condition and watcher	Per 9.4

Field	Purpose
Current state	Valid, expiring, expired

The retrieval standard: **if the artifact cannot be produced within one working day, it is not evidence.** Test this by sampling, not by asserting it.

9.6 Evidence architecture controls

BOARD EXTRACT, LIFTABLE

Our evidence position rests on three properties that are usually confused. Fidelity means the record describes what happened, which we achieve by having the component that authorizes an action also record it. Completeness means the record is not silently missing entries, which cryptography does not address and which we establish by sequence gaps, heartbeats and monthly reconciliation against our model provider's billing. Currency means the evidence is still true, which we achieve by attaching to every attestation the specific change events that void it and naming what watches for each. Evidence that expires by design is materially stronger than evidence that expires by audit discovery, because in the second case we are operating on false assurance in the interval.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
EVD-001	**Identify and document the enforcement point for each obligation requiring evidence, and record where no enforcement point exists**	An obligation with no enforcement point can only be narrated, and knowing that is better than discovering it	AI Gov Lead with Head of Engineering	Obligation-to-enforcement-point map, with gaps listed	EU: Arts. 12, 19 • NIST: MEASURE, MANAGE • ISO: 7.5, 8.1	ALL / HR / ALL	Once, then on architecture change	New system component, new integration, decommissioning of an enforcement point
EVD-002	**Configure the inference gateway to write the log record at the point the call is authorized, so no path reaches the model without emitting**	Evidence written downstream records what was reported, not what happened	Head of Engineering with CISO countersign	Gateway configuration under version control, plus bypass test report across all enumerated paths	EU: Arts. 12, 19 • NIST: MEASURE • ISO: 8.1, A.6 • SG: Agentic D3	ALL / HR / ALL	At deployment, then on any routing change	New endpoint, gateway version change, new direct-calling client, change of log destination
EVD-003	**Configure emission to fail closed, so that a system that cannot record cannot act, and document each exception as an accepted risk with a named accepter**	Fail-open logging produces its thinnest record during the incidents you will be asked about	Head of Engineering	Fail-closed configuration evidence, plus a register of exceptions with named acceptances	EU: Arts. 12, 15 • NIST: MANAGE • ISO: 8.1	ALL / HR / ALL	At deployment, then quarterly	New exception granted, availability requirement change, buffer or queue introduced

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
EVD-004	**Submit the write path design for governance review before implementation, covering call sites, synchronicity, failure behavior, sequencing and buffer overflow handling**	The completeness ceiling is set at write path design and cannot be raised downstream	Head of Engineering with AI Gov Lead	Write path design review record with governance sign-off	EU: Art. 12 • NIST: MEASURE • ISO: 8.1, 7.5	ALL / HR / ALL	Before implementation of any new emitter	New emitter, change to any of the five design decisions, new logging destination
EVD-005	**Emit a monotonic sequence per emitter and run automated gap detection, alerting on any discontinuity**	Without sequencing, missing records are indistinguishable from inactivity	Head of Engineering	Gap detection configuration and alert history including resolved gaps	EU: Art. 12 • NIST: MEASURE • ISO: 8.1	ALL / HR / ALL	At deployment, then continuous	New emitter, sequence reset, emitter replacement
EVD-006	**Emit periodic heartbeat records from every emitter so that silence is distinguishable from stopped recording**	An emitter producing nothing is ambiguous, and the ambiguity resolves against you	Head of Engineering	Heartbeat configuration and monitoring for heartbeat absence	EU: Art. 12 • NIST: MEASURE • ISO: 8.1	ALL / HR / ALL	At deployment, then continuous	New emitter, change to heartbeat interval, monitoring reconfiguration
EVD-007	**Reconcile the count of recorded inferences against the model provider's billing records at least monthly, and investigate any shortfall**	Provider billing is an independent count with no interest in your evidence position. It is the strongest completeness control available	Head of Engineering with Finance	Monthly reconciliation record showing counts, variance and investigation of shortfalls	EU: Arts. 12, 19 • NIST: MEASURE • ISO: 9.1	ALL / ALL / ALL	Monthly	New provider, change of billing metric, addition of a self-hosted model with no independent count
EVD-008	**Attach a validity condition to every control attestation, naming the voiding events, the watcher for each, and a hard expiry**	Evidence must expire by design. The alternative is discovering expiry during an audit	AI Gov Lead	Attestation records each carrying voiding events, watchers and expiry	EU: Art. 9(2) • NIST: MANAGE • ISO: 9.1, 10.3	ALL / ALL / ALL	At every attestation	Attestation made without a named watcher for any listed event
EVD-009	**Verify that each named watcher is actually implemented and firing, by inducing a test voiding event and confirming the attestation is marked expired**	A validity condition with an unimplemented watcher is an intention	AI Gov Lead with Engineering	Watcher test record per voiding event class	EU: Art. 9(2) • NIST: MANAGE • ISO: 9.1	ALL / HR / ALL	At implementation, then annually	New voiding event class, change to registry, pipeline or IAM tooling
EVD-010	**Maintain the evidence register mapping artifact to control to obligation to emitting component, and flag any artifact with no emitting component**	An artifact with no emitter was narrated. That is a finding you should raise on yourself	AI Gov Lead	Evidence register with emitter column populated and gaps flagged	EU: Art. 11, Annex IV • NIST: MEASURE • ISO: 7.5	ALL / ALL / ALL	Continuous	New artifact type, new obligation, change of emitter

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
EVD-011	**Sample five artifacts quarterly and time their production from request to delivery, recording any that exceed one working day**	Audit readiness is a measured property, not a claimed one	Internal Audit	Retrieval timing test records	EU: Art. 11 • NIST: MEASURE • ISO: 7.5	ALL / HR / ALL	Quarterly	Change of repository, change of retention tier, archival to cold storage
EVD-012	**Retain automatically generated logs for at least the period required, and record the retention basis per system**	Retention shorter than the obligation converts a compliant system into an evidential gap	Head of Engineering with Legal	Retention configuration per system with the basis recorded	EU: Arts. 19, 26 • NIST: MANAGE • ISO: 7.5	P, D / HR / EU	At deployment, then annually	Retention policy change, storage tier change, change to the applicable period

9.7 What this chapter changes downstream

Every other control chapter names an evidence artifact. This chapter determines whether those artifacts are worth anything. Three questions now attach to each one. Which enforcement point emitted it, or was it narrated? What establishes its completeness, or is completeness merely assumed? What voids it, and what is watching for that?

An artifact that cannot answer all three is a document. It may still be useful, and in many cases it is all that exists. But it should be recorded as a document, and the gap should be visible in the register rather than discovered by someone else.

CHAPTER 10

Human oversight and the defensible decision record

ENFORCEMENT POINT FOR THIS DOMAIN

The approval service. Oversight evidenced by screenshot or case note is narrated. Oversight where the action requires a token that only the approval service issues is enforced, and the token issuance is the record.

10.1 What Article 14 requires

High risk systems shall be designed and developed so they can be **effectively overseen by natural persons** during the period in use. Oversight measures must be commensurate with the risks, the level of autonomy, and the context of use. The operative words are effectively and by design. Oversight bolted on after build is the standard failure.

10.2 The four conditions of effective oversight

Condition	Test
Capability	Can the overseer understand the output well enough to disagree with it?
Authority	Can they override without seeking permission, and without career consequence?
Information	Do they see confidence, basis and limitation, not only the recommendation?
Time	Does the throughput target permit judgment?

A 100 percent approval rate is not evidence of a good model. It is evidence that the fourth condition has failed.

10.3 The defensible decision record

Element	Emitted by
The input the decision was made on	Inference gateway
The exact model and prompt version	Model registry, bound at inference
The output and its confidence	Inference gateway
Whether a human reviewed, and their decision	Approval service
The basis on which the human decided	Approval service, structured field, not free text

10.4 Human oversight controls

BOARD EXTRACT, LIFTABLE

The law requires that a person can effectively oversee each high risk system, meaning they can understand the output, disagree with it, and stop it. We treat oversight as measurable rather than declared. We track override rates and review times, and an approver at or near 100 percent agreement triggers an investigation, because that pattern indicates the person lacks the time, information or authority to exercise judgment. Approval is enforced technically: the action cannot execute without a token the approval service issues.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
HUM-001	**Produce an oversight design per high risk system naming the overseeing role, the intervention mechanism, and the information presented at decision time**	Oversight described in policy but not designed into the system cannot be exercised	Business owner with AI Gov Lead	Oversight design document per system, approved before release	EU: Art. 14(1) to 14(4) • NIST: GOVERN, MANAGE • ISO: 5.3, A.9 • SG: Agentic D2	ALL / HR / EU	Before deployment	Interface change, change of overseeing role, autonomy increase, new decision type

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
HUM-002	**Implement mandatory approval as a token the action requires, issued only by the approval service, so an unapproved action cannot execute**	Approval recorded beside the action can be skipped. Approval that gates the action cannot	Head of Engineering	Approval service configuration plus a test showing the action fails without a token	EU: Art. 14(4) • NIST: MANAGE • ISO: A.9 • SG: Agentic D2	ALL / HR / ALL	At deployment, then on change	New action type, new execution path, change to token validation, emergency bypass granted
HUM-003	**Capture the overseer's basis for decision as a structured field, not free text, and make it mandatory for overrides and approvals of adverse outcomes**	Free text is not analyzable, so oversight effectiveness cannot be measured	Business owner with Engineering	Structured decision records with basis populated	EU: Arts. 12, 14 • NIST: MEASURE • ISO: A.9, 7.5	ALL / HR / EU	Continuous	Change to the decision taxonomy, new decision type, interface change
HUM-004	**Measure and report override rate, disagreement rate and average review time per overseer per month, and investigate any overseer at or near 100 percent approval**	Oversight effectiveness is a measurable property. Rubber stamping is visible in the data	AI Gov Lead with Internal Audit	Monthly oversight effectiveness report with investigations	EU: Art. 14 • NIST: MEASURE • ISO: 9.1	ALL / HR / EU	Monthly	Change of overseer population, throughput target change, model performance change
HUM-005	**Set and enforce a maximum review throughput per overseer, derived from measured time-to-judgment, and record the derivation**	Time is the condition that fails silently. A target set for productivity defeats the control	Business owner with CHRO	Throughput limit with the measurement basis, and enforcement evidence	EU: Art. 14(4) • NIST: GOVERN • ISO: 7.1, A. 9	ALL / HR / EU	At design, then annually	Volume increase, staffing reduction, change to case complexity
HUM-006	**Implement countermeasures against automation bias, at minimum displaying model confidence and requiring a reason where the overseer agrees with a low-confidence output**	Article 14 names automation bias explicitly and it is rarely addressed in design	Business owner with Data Science	Interface evidence showing confidence display and reason prompts	EU: Art. 14(4) • NIST: MEASURE • ISO: A.9 • SG: Agentic D2	ALL / HR / EU	At design, then on interface change	Interface change, confidence calibration change, removal of any prompt

CHAPTER 11

Transparency, disclosure, labeling and AI literacy

ENFORCEMENT POINT FOR THIS DOMAIN

For disclosure, the rendering layer. For labeling, the content generation pipeline at the point of creation. For literacy, the learning management record, not the completion claim.

Both obligations in this chapter are live now. They were not deferred.

11.1 Article 50 transparency

Obligation	Applies to	Practical requirement
Disclose AI interaction	Systems interacting directly with people	Inform the person unless obvious from context
Mark synthetic content	Providers of systems generating synthetic audio, image, video or text	Machine-readable marking, detectable as artificially generated
Disclose deepfakes	Deployers generating or manipulating image, audio or video resembling real persons, places or events	Disclose that the content is artificially generated or manipulated
Disclose AI-generated text on matters of public interest	Deployers publishing such text	Disclose unless human review with editorial responsibility applies
Emotion recognition and biometric categorization notice	Deployers of such systems	Inform exposed persons

Exact paragraph allocation of Article 50(1) to (5) and the wording of the exemptions is on the verify list. Note the interaction with Chapter 3: an emotion recognition system that is lawful still carries a notice duty. One that is prohibited is not made lawful by disclosing it.

11.2 Article 4 AI literacy

Providers and deployers shall take measures to ensure, to their best extent, a **sufficient level of AI literacy** among their staff and other persons dealing with AI systems on their behalf, taking into account their technical knowledge, experience, education and training, and the context of use.

Three things follow. It is proportionate, so one generic module for all staff fails the taking into account test. It extends beyond employees to persons acting on your behalf, which reaches contractors and some outsourced operations. It is evidenced by records, so delivery without records is unevidenced compliance.

11.3 Transparency and literacy controls

BOARD EXTRACT, LIFTABLE

Two obligations already apply to us regardless of the high risk deferrals: telling people when they are interacting with AI or seeing AI-generated content, and ensuring staff who work with AI have literacy appropriate to their role. The literacy duty extends to contractors acting on our behalf and is evidenced by per-person records, not by having run a program.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
TRN-001	**Implement AI interaction disclosure in the rendering layer so it cannot be omitted by a client, and capture a rendered example per interface**	Disclosure implemented per client is disclosure that a new client omits	Product with Engineering	Rendered evidence per interface, dated, plus rendering layer configuration	EU: Art. 50(1) • NIST: MAP • ISO: A.8	ALL / LR, HR / EU	At release, then per interface change	New interface, new channel, client-side rendering change, redesign
TRN-002	**Mark synthetic content at the point of generation with machine-readable provenance, and test that the marking survives your own distribution pipeline**	Marking applied after generation is lost by any path that bypasses the marker	Product with Engineering	Provenance implementation record plus a survival test across each distribution path	EU: Art. 50(2) • NIST: MAP • ISO: A.8 • SG: Content provenance	P / ALL / EU, IN, CN	At release, then per pipeline change	New distribution channel, transcoding step added, new generation endpoint
TRN-003	**Define AI literacy requirements per role, differentiated by technical knowledge and context of use, and record the derivation**	Article 4 requires proportionality. A single module for all staff does not evidence it	CHRO with AI Gov Lead	Role-differentiated literacy matrix with rationale	EU: Art. 4 • NIST: GOVERN • ISO: 7.2	ALL / ALL / EU	Once, then annually	New role interacting with AI, new system category, change to context of use
TRN-004	**Deliver literacy training and hold per-person completion records covering employees and contractors acting on your behalf**	The obligation extends beyond employees, and records are the evidence	CHRO	Per-person completion records including contractors	EU: Art. 4 • NIST: GOVERN • ISO: 7.2, 7.3	ALL / ALL / EU	On joining, then on refresh cycle	New joiner, role change, contractor onboarding, refresh cycle elapsing

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
TRN-005	**Provide deployers with the information they need to meet their own obligations, and retain the pack issued per version**	A provider withholding this transfers risk it cannot lawfully transfer	Product with Legal	Deployer information pack, versioned and issued	EU: Arts. 13, 25 • NIST: MAP • ISO: A.8	P / HR / EU	At each release	Model change, limitation discovered, intended purpose change

CHAPTER 12

AI-specific security and adversarial robustness

ENFORCEMENT POINT FOR THIS DOMAIN

The guardrail and the tool broker. A guardrail evaluated in the application can be bypassed by any caller that skips the application. A guardrail in the path between caller and model cannot.

12.1 Why this is a separate chapter

Article 15 requires high risk systems to achieve appropriate levels of accuracy, robustness and cybersecurity, and to be resilient against attempts to alter their use, outputs or performance by exploiting vulnerabilities. That last clause is an AI-specific security obligation, and your existing security program was not designed against it.

The threats are enumerated in the OWASP Top 10 for LLM Applications and MITRE ATLAS. Neither is law. Both are the reference taxonomies your findings should be expressed in, so that they can be compared and tracked.

12.2 The security controls

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
SEC-001	**Place guardrails in the request path rather than the application, and test by calling the model endpoint directly to confirm the guardrail still applies**	A guardrail the caller can skip protects only compliant callers	CISO with Engineering	Path test report showing direct endpoint calls are still filtered	EU: Art. 15 • NIST: MEASURE, MANAGE • ISO: A.6	ALL / HR / ALL	At deployment, then quarterly	New endpoint, new client, guardrail version change, routing change
SEC-002	**Test for indirect prompt injection by planting adversarial instructions in content the system retrieves, not only in user input**	Indirect injection is the harder case and the one production RAG systems fail	CISO	Indirect injection test report with planted-content methodology	EU: Art. 15 • NIST: MEASURE • ISO: A.6 • OWASP LLM01	ALL / ALL / ALL	Before release, then quarterly	New retrieval source, new document type ingested, parser change

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
SEC-003	**Verify that guardrails fail closed under load and under dependency failure, by inducing both conditions**	A filter that fails open converts a traffic spike into a safety incident	CISO with Engineering	Fail-closed test report covering load and dependency failure	EU: Art. 15 • NIST: MEASURE • ISO: A.6	ALL / HR / ALL	At deployment, then annually	Capacity change, new guardrail dependency, timeout configuration change
SEC-004	**Maintain an AIBOM enumerating models, datasets, libraries, orchestration frameworks and tool servers, with versions pinned**	You cannot assess or patch what you have not enumerated	CISO with Engineering	Current AIBOM with pinned versions and generation date	EU: Art. 15 • NIST: MAP • ISO: A.6, A.10 • OWASP LLM03	ALL / ALL / ALL	At build, then per release	New dependency, framework upgrade, new tool server, model substitution
SEC-005	**Verify model artifact integrity by signature before load, and scan for unsafe deserialization**	Some model formats execute on load, which makes a download a code execution path	CISO	Signing verification and pre-load scan records	EU: Art. 15 • NIST: MEASURE • ISO: A.6 • OWASP LLM03	ALL / ALL / ALL	At every model load	New model source, format change, registry change
SEC-006	**Test tenant and user isolation by attempting cross-tenant retrieval as a low-privilege principal, and record the negative result**	Cross-tenant leakage is a breach, and isolation asserted in architecture is not isolation tested	CISO	Isolation test report showing failed cross-tenant attempts	EU: Art. 15 • NIST: MEASURE • ISO: A.6 • OWASP LLM08	ALL / ALL / ALL	At deployment, then quarterly	New tenant, index rebuild, permission model change, new retrieval path
SEC-007	**Express all AI security findings in OWASP LLM or MITRE ATLAS terms, and maintain a documented position on each OWASP LLM Top 10 item**	A shared taxonomy is what makes findings comparable across time and teams	CISO	OWASP LLM Top 10 position statement, plus findings mapped to ATLAS	EU: Art. 15 • NIST: MEASURE • ISO: A.6	ALL / ALL / ALL	Annually, and per assessment	New OWASP or ATLAS release, new system architecture, new threat class observed

CHAPTER 13

Third party, procurement and agent authorization

ENFORCEMENT POINT FOR THIS DOMAIN

For vendors, the procurement gate. For agents, the tool broker that authorizes each call. An agent framework that records its own actions is self-reported. The broker is the enforcement point.

13.1 Third party AI

Article 25 governs responsibilities along the value chain and is where role migration bites. A deployer becomes a provider, with full provider obligations, where it puts its name or trademark on a high risk system, makes a substantial modification, or changes the intended purpose such that the system becomes high risk. The largest unmanaged category is not procured AI. It is AI added to software already licensed, which arrives by release note.

13.2 Agent authorization, the question that matters

For agentic systems the governing question is narrow and answerable: **what may this agent initiate without a human in the path?**

Singapore's Model AI Governance Framework for Agentic AI is the most specific guidance available. Version 1.5, published 20 May 2026, updated 5 June 2026. It is structured around four dimensions: assess and bound the risks upfront; make humans meaningfully accountable; implement technical controls and processes; enable end-user responsibility. It describes eight agent components, distinguishes action space from autonomy as the two determining axes, and identifies three multi-agent patterns: sequential, supervisor and swarm.

It is voluntary. It creates no obligation. It is nonetheless the best available design reference, and the EU AI Act reaches agents only through Articles 14 and 15, drafted before agentic deployment was common.

Action class	Default posture
Read internal, non-sensitive	Autonomous

Action class	Default posture
Read sensitive or personal data	Autonomous within scoped permission, logged, sampled
Write internal, reversible	Autonomous, logged, reversible within a defined window
Write internal, irreversible	Human approval required
External communication in the organization's name	Human approval required
Financial commitment, contractual commitment	Human approval required, second approver above threshold
Change to the agent's own permissions, tools or memory	**Prohibited to the agent. Human-only path**

The last row is the one most often missed. An agent that can widen its own scope has no scope.

13.3 Third party and agent controls

BOARD EXTRACT, LIFTABLE

For every AI agent we run, we can state exactly what it may do without a human in the path. That boundary is enforced by the component that authorizes each tool call, not by the agent's own configuration. Agents cannot modify their own permissions, tools or memory. Every agent has its own identity tied to a named authorizing person, so any action can be attributed and any agent can be stopped individually. Containment is tested, including mid-task, and we record how long a full stop takes.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
TPY-001	**Require AI assessment before contract signature, enforced at the procurement gate, with no path to purchase that skips it**	Assessment after signature has no leverage	Procurement with AI Gov Lead	Gate configuration plus assessment records per contract	EU: Art. 25 • NIST: GOVERN, MAP • ISO: A.10	D / ALL / ALL	Every procurement	New purchasing route, renewal without reassessment, exception granted

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
TPY-002	**Obtain written vendor attestation of whether the product contains AI, whether it infers emotion, and whether customer data trains their models, refreshed at renewal**	Three questions that determine your exposure and that vendors will not volunteer	Procurement	Signed vendor attestations, versioned	EU: Arts. 5(1)(f), 25 • NIST: MAP • ISO: A. 10	D / ALL / EU	At procurement and renewal	Renewal, vendor product update, vendor acquisition
TPY-003	**Contract for model change notification with a defined notice period, and treat an unnotified model substitution as a contractual breach**	Vendor model swaps void your validation without any decision on your side	Procurement with Legal	Executed clauses plus a log of notifications received	EU: Arts. 25, 9(2) • NIST: MANAGE • ISO: A.10	D / HR / ALL	At contracting	Contract renewal, vendor architecture change, provider substitution by the vendor
AGT-001	**Issue each agent a distinct identity bound to an authorizing human, and prohibit shared service accounts for agent execution**	Without per-agent identity, actions cannot be attributed and cannot be revoked selectively	CISO	Agent identity register with authorization chain per agent	EU: Art. 14 • NIST: GOVERN • ISO: A.9 • SG: Agentic D2, D3	ALL / ALL / ALL	At agent creation	New agent, change of authorizing human, credential rotation, agent cloning
AGT-002	**Publish an authorization matrix per agent defining which action classes are autonomous and which require human approval, and configure the tool broker to enforce it**	Autonomy defined in a document and not in the broker is not a boundary	CISO with Engineering	Broker configuration under version control, matching the published matrix	EU: Arts. 14, 15 • NIST: MANAGE • ISO: A.9 • SG: Agentic D1, D2	ALL / ALL / ALL	At agent deployment	New tool granted, new action class, matrix revision, broker version change
AGT-003	**Prohibit agents from modifying their own permissions, tool set or memory, and route all such changes through a human-only path**	An agent that can widen its own scope has no scope	CISO with Engineering	Broker policy showing self-modification denied, plus a test attempt	EU: Art. 14 • NIST: MANAGE • ISO: A.9 • SG: Agentic D3	ALL / ALL / ALL	At deployment, then quarterly	New tool with configuration capability, framework upgrade, memory subsystem change
AGT-004	**Implement and test containment that halts an agent mid-task, recording the time from trigger to full stop**	Untested containment is an assumption. Mid-task is the case that matters	CISO	Containment test record with measured stop time	EU: Arts. 14, 15 • NIST: MANAGE • ISO: A.9 • SG: Agentic D3	ALL / ALL / ALL	At deployment, then quarterly	New execution environment, framework change, new long-running task type
AGT-005	**Log the full action sequence per agent task at the broker, not the framework, and review sampled trajectories monthly**	Harm from agents emerges across a chain of individually reasonable steps	CISO with Engineering	Broker-emitted trajectory logs plus monthly review records	EU: Art. 12 • NIST: MEASURE • ISO: A.6, A.9 • SG: Agentic D3	ALL / ALL / ALL	Continuous, review monthly	New tool, new agent pattern, change to log destination
AGT-006	**Assess multi-agent composition risk before deploying any sequential, supervisor or swarm pattern, and bound delegation depth**	Composition amplifies single-agent failure and is assessed at system level	CISO with Engineering	Composition risk assessment with delegation depth limit configured	EU: Arts. 9, 15 • NIST: MAP • ISO: 6.1, A.9 • SG: Agentic v1.5	ALL / ALL / ALL	Before deploying any multi-agent pattern	New agent added to a pattern, pattern change, third-party agent introduced

CHAPTER 14

Monitoring, incident response and serious incident reporting

ENFORCEMENT POINT FOR THIS DOMAIN

For monitoring, the metric pipeline with thresholds bound to actions. For incidents, the reporting clock, which starts on awareness rather than on confirmation.

14.1 Post-market monitoring

Article 72 requires providers to establish and document a post-market monitoring system proportionate to the nature of the AI technologies and the risks, actively and systematically collecting, documenting and analyzing relevant data on performance throughout the lifetime of the system. Whether the Commission template referred to in Article 72 has been adopted is on the verify list.

Monitoring is only a control where a threshold breach triggers a named action assigned to a named role. Dashboards without thresholds are observability.

14.2 Serious incident reporting

Article 73(1) requires providers of high risk systems placed on the Union market to report serious incidents to the market surveillance authorities of the Member States where the incident occurred. Serious incident is defined in **Article 3(49)**. The deadlines are tiered by severity and are set out below.

Deadline	Trigger	Provision
2 days	Widespread infringement, or a serious incident within Article 3(49)(b), being a serious and irreversible disruption of the management or operation of critical infrastructure	Art. 73(3)
10 days	Death of a person. Report immediately once a causal relationship is established or suspected	Art. 73(4)
15 days	All other serious incidents. Report immediately after establishing a causal link or the reasonable likelihood of one	Art. 73(2)

Three operational consequences. **The clock starts on awareness**, for the provider or, where applicable, the deployer, so your detection-to-escalation path sits inside the deadline rather than additional to it. **An incomplete initial report is expressly permitted**, to be followed by a complete report, which means the two-day deadline is achievable without full root cause. And the system **must not be altered** in a way that could affect the subsequent evaluation of causes before the authorities have been informed, so your containment plan and your evidence preservation plan must be reconciled in advance.

Deployers who identify a serious incident must inform the provider. Commission draft guidance issued 26 September 2025 addresses a simplified route where equivalent sector reporting already exists, for example critical infrastructure under NIS 2; whether that guidance is final is on the verify list. Note also that Article 55(1)(c) is a separate obligation for providers of GPAI models with systemic risk to inform the AI Office, and is not discharged by Article 73 reporting.

The decision to report is a legal judgment that must be pre-assigned to a named role, because a two-day deadline cannot be met by a committee convening.

14.3 Monitoring and incident controls

BOARD EXTRACT, LIFTABLE

Serious incidents involving high risk AI must be reported to regulators on a clock that starts when we become aware, not when we confirm. One named person, with a named deputy and out-of-hours cover, decides reportability. We test that path annually against the shortest applicable deadline and record the elapsed time. We preserve forensic material at declaration, because logs roll over on schedules that do not know an incident has occurred.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
MON-001	**Bind every monitored metric to a threshold and every threshold to a named action and role, and reject any metric without both**	A metric with no threshold and no owner is not a control	Engineering with CRO	Metric register showing threshold, action and role per metric	EU: Art. 72 • NIST: MEASURE • ISO: 9.1	ALL / HR / EU	At deployment, then quarterly	New metric, threshold revision, owner change, alert route change
MON-002	**Route alerts to a monitored channel with acknowledgment tracking and automatic escalation on non-acknowledgment**	An alert into an unwatched channel is not a control	Engineering with Ops	Alert routing configuration plus acknowledgment and escalation history	EU: Art. 72 • NIST: MANAGE • ISO: 9.1	ALL / HR / ALL	At deployment, then quarterly	Channel change, rota change, escalation policy change
MON-003	**Document the post-market monitoring plan per high risk system, stating what is collected, how it is analyzed, and how findings reach the Article 9 process**	Article 72 requires a system, not incidental telemetry	Engineering with CRO	Post-market monitoring plan plus collected data and analysis records	EU: Art. 72 • NIST: MEASURE • ISO: 9.1	P / HR / EU	Before market placement, then annually	Model change, new deployment context, template adoption by the Commission

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
INC-001	**Pre-assign, by name, the role that decides whether an incident is reportable, with a named deputy and out-of-hours coverage**	Reportability is a legal judgment on a clock. It cannot be made by committee	General Counsel	Named decision authority with deputy, tested out of hours	EU: Art. 73 • NIST: MANAGE • ISO: 10.2	P / HR / EU	Once, then on personnel change	Departure or role change of primary or deputy, change of jurisdiction
INC-002	**Build a notification matrix stating, per jurisdiction and incident type, the recipient authority, the deadline and the clock start event, and encode the EU tiers of 2, 10 and 15 days against their Article 3(49) triggers**	The EU two-day tier is unachievable without a pre-built matrix, and regimes run concurrently	General Counsel with AI Gov Lead	Notification matrix carrying the Art. 73(2), (3) and (4) deadlines, reviewed quarterly	EU: Arts. 3(49), 73(2) to 73(4) • NIST: MANAGE • ISO: 10.2	P, D / HR / ALL	Quarterly	New jurisdiction, regulatory change, new incident category
INC-003	**Preserve forensic material on incident declaration, including prompts, outputs, model and prompt versions, agent trajectories and gateway logs, before any retention period expires**	Evidence rolls over on a schedule that does not know an incident occurred	CISO with Engineering	Preservation procedure plus per-incident preservation record	EU: Arts. 12, 19, 73 • NIST: MANAGE • ISO: 10.2	ALL / ALL / ALL	On every incident declaration	Retention change, new log source, storage tier change
INC-004	**Run a tabletop exercise at least annually against an AI-specific scenario, timing the detection-to-reporting-decision path against the two-day Article 73(3) deadline as the binding case**	Two days is the shortest clock. A process that clears fifteen days may still fail the case that matters	CISO with General Counsel	Exercise record with measured elapsed time against deadline	EU: Art. 73 • NIST: MANAGE • ISO: 10.2	P / HR / EU	Annually	Deadline change, personnel change, new system category, prior exercise failure
INC-005	**Log near misses to the same register as incidents, with a defined near-miss criterion**	Near misses are the cheapest available evidence of where controls are thin	CISO	Incident register including near misses with criterion applied	EU: Art. 72 • NIST: MANAGE • ISO: 10.2	ALL / ALL / ALL	Continuous	Criterion revision, register migration

CHAPTER 15

Maturity model and the board pack

Ten dimensions, five levels. Level 3 is the target, because it is the point at which a program becomes defensible. Levels 4 and 5 are differentiators.

15.1 The maturity model

Level	Name	Definition
1	Initial	AI in use without systematic governance. Cannot answer how many systems exist
2	Developing	Partial coverage. Inventory started, policy exists, high-profile systems assessed
3	**Defined**	Complete inventory, consistent assessment, named ownership, working gate, evidence repository. **Defensible**
4	Managed	Measured. Controls validated for effectiveness, monitoring continuous, board reporting meaningful
5	Optimized	Automated evidence, anticipates regulation, contributes to standards

#	Dimension	Level 3 test
1	Strategy and leadership	Board-approved risk appetite, named accountable executive, funded program
2	Policy framework	Approved, enforced, reviewed on cycle
3	Risk management	Assessment per system, refreshed on defined triggers
4	Control implementation	One control set mapped to all instruments, recorded in a Statement of Applicability

#	Dimension	Level 3 test
5	Technology and tooling	Enforcement points emit; gate cannot be bypassed
6	Evidence and documentation	Artifact-to-control index current, retrieval within one working day
7	Monitoring and post-market	Thresholds bound to actions, breaches reach the risk process
8	Assurance and audit	Internal audit operating, effectiveness tested not existence
9	Third party and supply chain	Assessment before contract, embedded AI screened
10	People and AI literacy	Role-differentiated, delivered, recorded, refreshed

15.2 The board pack

Four numbers, one judgment, one page. Everything else is appendix.

Metric	Source	Healthy
Systems inventoried and classified	Inventory register	Approaching 100 percent, sweep current
P1 controls evidenced, of 83	Evidence register	Rising; 100 percent before any high risk deadline
Maturity level	Chapter 15 assessment	3.0 or above
Gaps with owner and date	Gap report	Every gap has both. Zero unowned

Three risk indicators, reported every meeting: unclassified systems in production, overdue reassessments, open serious incidents or reportable events.

The judgment, one sentence: whether the program is defensible today, and if not, the date it will be.

BOARD EXTRACT, LIFTABLE

This entire section is the board pack. Report the four metrics and three risk indicators at every meeting, the maturity score quarterly, and a full governance review annually or on material change. Do not present 260 controls to a board. Present four numbers, three indicators, and a judgment.

CHAPTER 16

Phased implementation pathway

Sequence matters because the P1 controls are not independent. Classification requires inventory. Assessment requires classification. Evidence mapping requires a control set.

Wave	Days	Do	Exit test
1. Visibility	1 to 30	Appoint the accountable lead with budget. Discovery sweep including shadow AI and embedded SaaS features. Inventory with a named owner per system. Article 5 screening across the whole inventory	You can state what AI you run, who owns each system, and that nothing prohibited is running
2. Classification and framework	31 to 60	Article 2 applicability memo. Role per system. Tier per system with rationale. Article 6(3) filter documented where relied on. Adopt the crosswalk as the single control set. Define the evidence model	Every system classified with written rationale. One control framework adopted
3. Enforcement points	61 to 90	Gateway emission at authorization. Deployment gate that cannot be bypassed. Approval tokens for mandatory human decisions. Sequence numbering and heartbeats. First provider reconciliation	Evidence is emitted rather than narrated, and you have tested that no path bypasses the gate
4. High risk core	91 to 180	Risk and impact assessments. Oversight design and effectiveness measurement. Data provenance and retrieval permission enforcement. Incident path with named reportability authority, tested	High risk systems have current assessments, working oversight, complete logs, and a tested reporting path
5. Assurance	180+	Monitoring with bound thresholds. Adversarial testing. Vendor assessments. First internal audit. First board pack. Conformity preparation ahead of the applicable date	Level 3 maturity. Independently assured

Role cuts for distribution

CISO: Chapters 12, 13 agent controls, 14 incident, roughly 50 controls. **CRO:** Chapter 7, impact assessments, roughly 30. **Legal and compliance:** Chapters 2, 3, 11, 14 reporting, roughly 30. **Engineering:** Chapters 9, 12 technical, monitoring implementation, roughly 60. **CDO:** Chapter 8. **HR:** Chapter 11 literacy. **Procurement:** Chapter 13. **Internal audit:** assurance only, and nothing else, because independence is the point.

CHAPTER 17

Policy library

Fifteen policies. Each is a specification, not a template. Write them in your own voice and against your own operating model.

Cross-cutting drafting rules. Every policy names one accountable owner, not a function. Every policy states what evidence its operation generates, and that evidence appears in the register from Chapter 9. Every policy carries a review trigger that is event-based as well as periodic.

1. Enterprise AI Acceptable Use Policy

Govern how staff may use AI tools. Applies to all employees and contractors, all AI tools, approved or otherwise.

MANDATORY CLAUSES Approved tool list and how to request additions; data classification rules for prompts, stating what must never be entered; prohibition on entering personal, confidential or regulated data into unapproved tools; mandatory verification of output before reliance; disclosure expectations when AI-assisted output is published or sent externally; consequences of breach.

EVIDENCE GENERATED Acknowledgment records per person; approved tool register; exception log

Owner: AI Governance Lead • **Mapping:** EU Art. 4 • ISO A.9, 7.3 • NIST GOVERN • **Review trigger:** New tool category, incident involving misuse, annual

2. AI System Inventory and Registration Policy

Ensure every AI system is known, classified and owned. Covers all AI systems including embedded, pilot and shadow.

MANDATORY CLAUSES Definition of what constitutes an AI system for registration; mandatory registration before deployment, enforced at the gate; required fields per Chapter 6; named accountable individual per system; quarterly SaaS embedded-AI screening; EU database registration where applicable.

EVIDENCE GENERATED Inventory register; gate configuration; sweep reports

Owner: AI Governance Lead • **Mapping:** EU Arts. 16, 49, 71 • ISO 4.3 • NIST MAP • **Review trigger:** New deployment route, acquisition, annual

3. AI Risk Management Policy

Define how AI risk is assessed, treated and accepted. All systems; full regime for high risk.

MANDATORY CLAUSES Assessment methodology and scoring scales; classification procedure and rationale requirement; the full reassessment trigger list; risk acceptance thresholds and who may accept at each level; requirement that acceptance is named, dated and conditional; integration with enterprise risk.

EVIDENCE GENERATED Methodology; risk register; per-system assessments; signed acceptances

Owner: Chief Risk Officer • **Mapping:** EU Art. 9 • ISO 6.1 • ISO/IEC 23894 • NIST MAP, MEASURE • **Review trigger:** Appetite change, regulatory change, annual

4. Prohibited Practices Policy

State what the organization will not build or deploy, irrespective of commercial case. All AI, all tiers, all jurisdictions.

MANDATORY CLAUSES The eight Article 5 practices restated as internal prohibitions; explicit prohibition on inferring emotional state of employees, candidates or students, with the medical and safety exceptions defined narrowly and requiring counsel sign-off; mandatory pre-deployment screening of 100 percent of inventory; vendor attestation requirement; escalation route where a proposed use approaches a line; no commercial override.

EVIDENCE GENERATED Screening register; counsel sign-offs; vendor attestations

Owner: General Counsel • **Mapping:** EU Art. 5 • ISO 6.1.2, A.5 • NIST MAP, GOVERN • **Review trigger:** Amendment to Article 5, new guidance, new HR or edtech tool, annual

5. Data Governance and Provenance Policy

Govern data used to train, tune, evaluate and run AI, including derived artifacts.

MANDATORY CLAUSES Eligibility criteria; provenance and original collection purpose recording; dataset versioning bound to model versions; preparation operation logging; bias examination and mitigation; data gap recording; retrieval permission enforcement; retention covering embeddings, indexes, caches and logs; special category processing for bias correction with safeguards.

EVIDENCE GENERATED Lineage records; version register; bias reports; gap register; retrieval permission tests

Owner: Chief Data Officer • **Mapping:** EU Art. 10, GDPR • ISO A.7 • NIST MAP, MEASURE • **Review trigger:** New data source, new retrieval architecture, annual

6. Human Oversight Policy

Define where humans must be able to intervene and with what authority. All high risk and all consequential automated decisions.

MANDATORY CLAUSES Decisions requiring oversight and at what level; overseer authority to override, halt and escalate without permission; mandatory approval conditions; approval enforced by token; structured decision-basis capture; automation bias countermeasures; maximum throughput derived from measured judgment time; effectiveness measurement including override rate.

EVIDENCE GENERATED Oversight designs; approval logs; effectiveness reports

Owner: AI Governance Lead with business owners • **Mapping:** EU Art. 14 • ISO A.9, 5.3 • NIST GOVERN, MANAGE • SG Agentic D2 • **Review trigger:** New decision type, effectiveness metric breach, annual

7. Agentic and Autonomous Action Policy

Define what an agent may initiate without a human in the path. All agentic systems, internal and third party.

MANDATORY CLAUSES Autonomy levels and permitted action space; the authorization matrix per action class; per-agent identity bound to an authorizing human; prohibition on shared service accounts; least-privilege tool scoping; prohibition on agent self-modification of permissions, tools or memory; containment and kill switch with testing cadence; delegation depth limits for multi-agent patterns; broker-emitted trajectory logging.

EVIDENCE GENERATED Identity register; broker configuration; containment tests; trajectory logs

Owner: CISO with CTO • **Mapping:** EU Arts. 14, 15 • ISO A.9, A.6 • NIST GOVERN, MANAGE • SG Agentic D1 to D4 • **Review trigger:** New agent pattern, new tool class, framework change, semi-annual

8. Model Change and Revalidation Policy

Ensure change does not silently void validation. All model, prompt, tool and configuration changes, including vendor-initiated.

MANDATORY CLAUSES What constitutes a material change; mandatory revalidation triggers including vendor model substitution; prompt changes treated as changes; rollback requirement; version binding of model, prompt and dataset; contractual notice requirement from vendors; release gate tied to governance approval.

EVIDENCE GENERATED Change records; revalidation results; version register; vendor notifications

Owner: CTO • **Mapping:** EU Arts. 9(2), 15, 17, 25 • ISO 8.1, ISO/IEC 5338 • NIST MANAGE • **Review trigger:** New change type, vendor architecture change, annual

9. Evaluation and Testing Policy

Define what testing is required before and after release. All high risk and generative systems.

MANDATORY CLAUSES Acceptance criteria defined before build; testing against pre-agreed thresholds; adversarial testing including indirect prompt injection; guardrail bypass and fail-closed testing; bias testing across groups defined by risk not data availability; isolation testing; cadence and re-test triggers; results retained as evidence.

EVIDENCE GENERATED Test plans; results dated before release; red team reports

Owner: Head of Engineering with QA • **Mapping:** EU Arts. 9, 15 • ISO 8.1, A.6 • NIST MEASURE • SG Testing and assurance • **Review trigger:** New threat class, new system category, annual

10. Logging and Record Retention Policy

Establish what is recorded, by what, and for how long. All AI systems; full requirements for high risk.

MANDATORY CLAUSES Emission at the enforcement point, not downstream; fail-closed emission with documented exceptions; mandatory record fields; sequence numbering and heartbeats; monthly reconciliation against provider billing; retention periods per obligation; preservation on incident; write path design review before implementation.

EVIDENCE GENERATED Gateway configuration; bypass tests; gap detection; reconciliation records

Owner: Head of Engineering with Legal • **Mapping:** EU Arts. 12, 19, 26 • ISO 7.5, 8.1 • NIST MEASURE • SG Agentic D3 • **Review trigger:** New emitter, retention change, new obligation, annual

11. Transparency and Disclosure Policy

Govern what is told to users, affected persons and downstream deployers. All user-facing and content-generating systems.

MANDATORY CLAUSES AI interaction disclosure implemented in the rendering layer; synthetic content marking at generation with provenance metadata; deepfake disclosure; documented limitations and intended use; deployer information packs; information to affected persons including how to seek human review; retention of disclosure evidence.

EVIDENCE GENERATED Rendered disclosure examples; provenance implementation and survival tests; deployer packs

Owner: Product with Legal • **Mapping:** EU Arts. 13, 50, 86 • ISO A.8, ISO/IEC 12792 • NIST MAP • **Review trigger:** New channel, new content type, guidance change, annual

12. AI Literacy and Training Policy

Ensure people working with AI are competent to do so. Employees and contractors acting on the organization's behalf.

MANDATORY CLAUSES Role-differentiated literacy requirements with derivation; delivery and per-person recording; contractor coverage; oversight-specific training; board education; refresh cycle; onboarding inclusion; effectiveness assessment.

EVIDENCE GENERATED Literacy matrix; per-person completion records; competence matrix

Owner: CHRO with AI Governance Lead • **Mapping:** EU Art. 4 • ISO 7.2, 7.3 • NIST GOVERN • **Review trigger:** New role category, new system type, annual

13. AI Incident Response and Regulatory Reporting Policy

Convert a reporting obligation with a clock into a process that runs. All AI incidents and near misses.

MANDATORY CLAUSES AI-specific severity classification; detection and triage; named reportability decision authority with deputy and out-of-hours cover; notification matrix by jurisdiction with clock start defined as awareness; forensic preservation on declaration; containment and rollback authority; post-incident review; near-miss capture; annual tested exercise.

EVIDENCE GENERATED Incident register including near misses; notification records; exercise reports with timings

Owner: CISO with General Counsel • **Mapping:** EU Art. 73 • ISO 10.2 • NIST MANAGE • SG Incident reporting • **Review trigger:** Deadline change, new jurisdiction, post-incident, semi-annual test

14. AI Procurement Policy

Prevent AI risk entering through the buying process. All purchases, renewals and free-tier adoptions.

MANDATORY CLAUSES Assessment before signature, enforced at the procurement gate; vendor attestation on AI presence, emotion inference and training on customer data; mandatory contract clauses covering model change notice, audit rights, incident notification, data use and deletion; provenance verification; provider versus deployer role determination; quarterly embedded-AI screening of existing SaaS; offboarding and deletion verification.

EVIDENCE GENERATED Assessments; executed clauses; attestations; screening records

Owner: Procurement with Legal • **Mapping:** EU Arts. 25, 26 • ISO A.10 • NIST GOVERN, MAP • **Review trigger:** New purchasing route, vendor incident, annual

15. AI Intellectual Property and Confidentiality Policy

Protect the organization's information and clarify ownership of AI output. All AI use touching proprietary, client or licensed material.

MANDATORY CLAUSES What may not be entered into third-party models; ownership position on AI-generated output; contractual position on vendor training rights over your data; handling of client material under confidentiality obligations; licensed and open source data use in training, and license compatibility; provenance requirements where output is published; treatment of model weights and prompts as protectable assets.

EVIDENCE GENERATED License register; contractual positions; training data license records

Owner: General Counsel with CTO • **Mapping:** EU Art. 53 • ISO A.7, A.10 • NIST GOVERN • **Review trigger:** New model provider, litigation development, copyright guidance change, annual

CHAPTER 18

ISO/IEC 42001 certification readiness

ISO/IEC standards are copyrighted commercial publications. This chapter reproduces no clause text, no requirement wording, no Annex A control titles and no annex content. Clause numbers appear only as citations.

18.1 How this checklist is built, and why

That constraint produced a better instrument. A checklist that transcribes the standard tells you what the standard says, which you can learn faster by buying it. This chapter tells you something the standard does not: **what an accredited auditor will ask you to put on the table, and whether you can produce it.**

So the organizing principle is not the clause structure. It is the audit itself. Every row names an artifact, states what it must demonstrate when someone hostile to your interests examines it, and points to where in this guide it was produced.

You will still need a licensed copy of ISO/IEC 42001 to implement against. This chapter prepares you for the audit. It does not substitute for the standard, and any document that claims to is either infringing or wrong.

18.2 What certification actually assesses

Certification is issued by a certification body, whose own competence requirements are set out in ISO/IEC 42006. Check accreditation before you engage: a certificate from an unaccredited body is a document, not assurance.

	Stage 1	Stage 2
Tests	Does the management system exist, is it documented, is it internally coherent, are you ready to be audited	Does it operate, is it effective, has the cycle turned

	Stage 1	Stage 2
Fails on	Missing documented information, a Statement of Applicability that does not reconcile to your risk treatment, scope that does not match reality	Records that show the process was designed but never ran
Typical outcome	Delay, re-scheduling	Nonconformities requiring corrective action and re-audit

The single most common Stage 2 failure is chronological.

Organizations arrive with excellent documentation and no evidence that the cycle has turned even once. Internal audit has not run. Management review has not happened. There are no nonconformities recorded, which auditors read not as excellence but as a system that is not operating. You need elapsed operating time before Stage 2, not just completed documentation.

18.3 Stage 1 readiness: can you evidence that the system exists

Ref	What the auditor will ask you to produce	The test it must pass	Where produced	A / C / E
R1-01	A written scope statement for the management system	It matches what you actually run. Auditors compare it against your AI inventory and probe exclusions	Ch 2, Ch 6	☐ A ☐ C ☐ E
R1-02	Identification of the internal and external issues and interested parties bearing on your AI	It is specific to you. Generic text signals a document written to pass rather than to be used	Ch 2	☐ A ☐ C ☐ E

Ref	What the auditor will ask you to produce	The test it must pass	Where produced	A / C / E
R1-03	An approved AI policy	Approved at the right level, dated, versioned, consistent with the other policies in Ch 17	Ch 17, policy 1	☐ A ☐ C ☐ E
R1-04	Evidence of top management commitment	Resourcing, attendance and decisions, not a signed foreword	Ch 6, Ch 17	☐ A ☐ C ☐ E
R1-05	Assignment of roles, responsibilities and authorities	One accountable individual per activity. Auditors test this by asking a named person what they own	Ch 6 INV-004	☐ A ☐ C ☐ E
R1-06	Your AI risk assessment method, documented	Repeatable by someone else. Two assessors using it should reach comparable results	Ch 7 RSK-001	☐ A ☐ C ☐ E
R1-07	Completed risk assessments	Per system, dated, with named residual risk acceptance	Ch 7 RSK-002, RSK-004	☐ A ☐ C ☐ E
R1-08	Your AI impact assessment method and completed assessments	Distinct from risk assessment. It addresses effects on people, not effects on you	Ch 7	☐ A ☐ C ☐ E
R1-09	**The Statement of Applicability**	See 18.4. This is where Stage 1 most often stalls	Ch 5, Ch 7	☐ A ☐ C ☐ E
R1-10	Objectives for the management system, with plans to achieve them	Measurable, owned, resourced, dated	Ch 15, Ch 16	☐ A ☐ C ☐ E
R1-11	Competence requirements and records for people whose work affects AI outcomes	Role-differentiated. A single all-staff module does not evidence competence	Ch 11 TRN-003, TRN-004	☐ A ☐ C ☐ E
R1-12	Evidence that awareness has been created, distinct from training delivered	People can describe the policy and their part in it. Auditors test by asking them	Ch 11, Ch 17	☐ A ☐ C ☐ E

Ref	What the auditor will ask you to produce	The test it must pass	Where produced	A / C / E
R1-13	Your control over documented information	Version control, access control, retention, and a demonstrable current version	Ch 9 EVD-010, EVD-012	☐ A ☐ C ☐ E
R1-14	Operational planning and control for the AI lifecycle	Gates exist, are named, and are enforced rather than described	Ch 6 INV-002, Ch 7 RSK-002	☐ A ☐ C ☐ E
R1-15	Your approach to third parties and suppliers in the AI value chain	Assessment before contract, and contractual controls actually executed	Ch 13, Ch 17 policy 14	☐ A ☐ C ☐ E

18.4 The Statement of Applicability

The SoA is the single artifact most likely to delay certification. It is the one document that must reconcile three things simultaneously: your risk assessment, your control decisions, and what you actually operate. Auditors know this, so they test it hard.

What it must do. For every control in the standard's reference set, record whether you apply it, and why. Inclusions must trace to a risk. Exclusions must carry a justification that survives challenge.

The justification is a restatement. Not applicable because we do not do this is a conclusion, not a justification. What survives is specific and factual: we do not develop models, we deploy vendor models under contract, unmodified, with no fine-tuning, and provider-side controls are addressed through contractual obligation.

The SoA does not reconcile to the risk treatment. A control is marked applicable but appears nowhere in the risk treatment plan, or a treated risk relies on a control marked excluded. Auditors trace both directions.

It is stale. Written for Stage 1, never revised, and by Stage 2 it describes controls that have changed. Give the SoA a validity condition, per Chapter 9, with the model registry and control framework as watchers.

Exclusions are made for convenience. An exclusion because implementation is hard is not a justification. An exclusion because the control addresses a risk you demonstrably do not carry is.

18.5 Stage 2 readiness: can you evidence that it operates

Ref	What the auditor will ask you to produce	The test it must pass	Where produced	A / C / E
R2-01	Records showing the operational controls ran, not merely that they exist	Records emitted by the enforcement point rather than narrated	Ch 9 EVD-002	☐ A ☐ C ☐ E
R2-02	Monitoring and measurement results, with what you did about them	A threshold breach with no consequent action is worse than no threshold	Ch 14 MON-001	☐ A ☐ C ☐ E
R2-03	**At least one completed internal audit** of the management system	Conducted by someone independent of what they audited. Auditors check the independence	Ch 15, Ch 16	☐ A ☐ C ☐ E
R2-04	Internal audit findings, and evidence they were tracked to closure	Findings closed with verification that the fix worked, not merely that it was applied	Ch 16	☐ A ☐ C ☐ E
R2-05	**At least one completed management review**, minuted, with decisions and actions	Decisions, not attendance. A review that produced no decisions did not review anything	Ch 15	☐ A ☐ C ☐ E
R2-06	Recorded nonconformities and the corrective action taken	An empty nonconformity register is a red flag, not a clean record	Ch 16	☐ A ☐ C ☐ E
R2-07	Evidence that corrective action addressed causes, not only symptoms	The record shows cause analysis, and whether the same nonconformity recurred	Ch 16	☐ A ☐ C ☐ E
R2-08	Evidence of continual improvement over a period	The system changed in response to what it learned. This needs elapsed time	Ch 15, Ch 16	☐ A ☐ C ☐ E

Ref	What the auditor will ask you to produce	The test it must pass	Where produced	A / C / E
R2-09	Competence and awareness evidenced at the individual level	Auditors interview people. Records that do not match what people say is the fastest route to a finding	Ch 11 TRN-004	☐ A ☐ C ☐ E
R2-10	Evidence that documented information is current and controlled in practice	Sampled. A superseded version in active use is a finding	Ch 9 EVD-010	☐ A ☐ C ☐ E
R2-11	Evidence of lifecycle gates operating, including at least one refusal or exception	A gate that has never stopped anything invites the question of whether it can	Ch 6 INV-002	☐ A ☐ C ☐ E
R2-12	Evidence of third-party controls operating, not just contracted	Assessments completed, notifications received and acted on	Ch 13 TPY-001 to TPY-003	☐ A ☐ C ☐ E
R2-13	Incident records, including near misses, with what changed as a result	An organization with no AI incidents in twelve months is either very small or not detecting	Ch 14 INC-005	☐ A ☐ C ☐ E
R2-14	A current Statement of Applicability that still reconciles	Re-tested at Stage 2 against what you now operate	18.4	☐ A ☐ C ☐ E
R2-15	Evidence that the impact assessment process ran on a real system and changed something	An assessment that has never altered a design reads as ceremony	Ch 7	☐ A ☐ C ☐ E

18.6 Scoring and what it means

Position	Reading
Any Stage 1 row Absent	Do not book Stage 1
All Stage 1 Evidenced, R2-03 or R2-05 Absent	Stage 1 is feasible. Stage 2 is not, regardless of everything else
More than three Stage 2 rows Claimed	You will collect nonconformities. Convert them before booking
All Evidenced, cycle turned at least once	Book it

> **On the relationship between certification and compliance.**
>
> A certificate demonstrates that you operate a management system that meets the standard. It is not evidence of compliance with the EU AI Act, and no certification body can issue one that is. Do not let a certificate be presented internally as legal assurance.

18.7 Readiness controls

BOARD EXTRACT, LIFTABLE

ISO/IEC 42001 certification is assessed in two stages. The first tests whether the management system exists and is documented. The second tests whether it operates, and that is where most organizations fail, because it requires evidence that the cycle has already turned at least once: an internal audit completed, a management review held, nonconformities recorded and corrected. That evidence cannot be produced at short notice, so certification timing is set by elapsed operating time rather than by documentation effort. A certificate demonstrates a functioning management system. It is not, and cannot be, evidence of compliance with the EU AI Act.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
ISO-001	**Verify the certification body's accreditation status before engagement, and record the accreditation body and scope**	An unaccredited certificate is a document, and customers increasingly check	Compliance	Accreditation verification record	ISO/IEC 42006	ALL / ALL / ALL	Before engagement, then at renewal	Change of certification body, lapse of their accreditation, scope extension
ISO-002	**Build the Statement of Applicability against a licensed copy of the standard, with every inclusion traced to a risk and every exclusion carrying a factual justification**	The SoA is the artifact most likely to stall Stage 1, and reconciliation is what it is tested on	AI Governance Lead	Current SoA reconciling to the risk treatment plan in both directions	ISO: 6.1.3 • NIST: GOVERN	ALL / ALL / ALL	Before Stage 1, then on control change	New or retired control, change of scope, change to risk treatment, model or vendor change
ISO-003	**Schedule the first internal audit and first management review to complete far enough before Stage 2 that their outputs and resulting corrective actions are themselves evidenced**	Stage 2 fails most often because the cycle has not turned, which is a scheduling failure	AI Governance Lead	Audit and review schedule worked backward from the target certification date	ISO: 9.2, 9.3	ALL / ALL / ALL	At program start	Target date change, audit or review slippage, certification body scheduling change
ISO-004	**Conduct the internal audit using auditors independent of the activity audited, and record the basis of their independence**	Independence is tested directly, and a self-audited process is a finding regardless of quality	Internal Audit	Audit report naming auditors and their independence basis	ISO: 9.2 • EU: Art. 17 • NIST: GOVERN	ALL / ALL / ALL	At least annually	Organizational change affecting independence, change of audit resource
ISO-005	**Record nonconformities as they arise rather than at audit time, with cause analysis and verification that the corrective action worked**	An empty nonconformity register reads as a system that is not operating	AI Governance Lead	Nonconformity register with cause analysis and closure verification	ISO: 10.2 • EU: Art. 72 • NIST: MANAGE	ALL / ALL / ALL	Continuous	Register migration, change of nonconformity criterion
ISO-006	**Attach a validity condition to the Statement of Applicability naming the change events that void it, with the control framework and model registry as watchers**	An SoA written for Stage 1 and unrevised is the second most common finding at Stage 2	AI Governance Lead	SoA carrying voiding events and named watchers	ISO: 6.1.3, 9.1 • Ch 9 EVD-008	ALL / ALL / ALL	At SoA issue	Any listed voiding event, hard expiry, scope change

CHAPTER 19

General purpose AI and foundation models

Most enterprises now build on models they did not train. This chapter covers what that means legally, and the specific ways an organization becomes a GPAI provider without deciding to.

ENFORCEMENT POINT FOR THIS DOMAIN

The model registry and the procurement gate. Whether you hold GPAI obligations turns on facts about the model and what you did to it, both of which are recorded at these two points or nowhere.

19.1 What a GPAI model is, and the two tiers

Article 3(63) defines a general purpose AI model as one trained on broad data, displaying significant generality, capable of competently performing a wide range of distinct tasks, and capable of being integrated into a variety of downstream systems. A model trained for a single narrow task is not GPAI. The Commission's GPAI guidelines treat training compute above 10 to the power 23 floating point operations, combined with generative capability across language, image or video, as the indicative threshold, while excluding specialized models that lack breadth.

There are two obligation tiers, and the second is an addition to the first rather than a replacement.

Tier	Trigger	Obligations
All GPAI models	Meeting the Article 3(63) definition	Article 53 baseline, plus Article 54 where the provider is established outside the Union
GPAI with systemic risk	Presumed where cumulative training compute exceeds 10 to the power 25 FLOPs (Art. 51(2)), or by Commission designation on Annex XIII criteria	Article 53 **and** Article 55

GPAI obligations have been enforceable since **2 August 2025**. Enforcement sits with the Commission and the AI Office under **Article 101**, separately from the national market surveillance route and the Article 99 penalty scheme.

19.2 The Article 53 baseline

Obligation	What it means in practice
Technical documentation of the model	Covering the training and testing process and the results of evaluation. Maintained and kept current, and provided to the AI Office or national authorities on request
Information to downstream providers	Sufficient for them to understand capabilities and limitations and to meet their own obligations. If you integrate a model, this is the pack you should be demanding from your provider
A copyright policy	A policy for compliance with Union copyright law, including the reservation of rights. It may be a single policy applying across all your models
A public summary of training content	Published using the template issued by the AI Office

The open source carve-out is narrower than it sounds. Providers releasing models under a free and open license are relieved of parts of the baseline, retaining the copyright policy and the training content summary. The relief **does not apply** where the model presents systemic risk.

19.3 Systemic risk, and the notification clock

Where the compute presumption is met, or is foreseeably about to be met, **Article 52 requires notification to the Commission without delay and in any event within two weeks**. The notification may include a substantiated argument that the model exceptionally does not present systemic risk. Reaching the threshold without notifying is itself a breach.

Article 55 then adds four obligations: model evaluation to standardised state of the art protocols including documented adversarial testing; assessment and mitigation of systemic risks at Union level; tracking, documenting and reporting serious incidents and corrective measures to the AI Office without undue delay; and an adequate level of cybersecurity for both the model and its physical infrastructure.

The General Purpose AI Code of Practice, published by the Commission on 10 July 2025, is voluntary and offers a route to demonstrate compliance with Articles 53 and 55. Signing it is not required and not signing it is not a breach.

19.4 How an enterprise becomes a GPAI provider by accident

This is the section most deployers need. You do not have to train a frontier model to acquire provider obligations.

What you did	Likely consequence
Called a hosted model through an API, unmodified	Deployer of an AI system. No GPAI provider obligations. Demand the Article 53 downstream pack from your provider
Fine-tuned a third-party model and kept it internal	Assess carefully. Modification can transfer obligations for the modified model. The Commission guidelines address lifecycle and post-market modification
Fine-tuned a model and made it available to others, including group entities under your own name	Strong indication you are placing a GPAI model on the market as its provider
Distilled, merged or continued pre-training on a base model and released the result	Treat as provider unless advised otherwise
Wrapped a model in your product and branded it as your own	Article 25 role migration for the AI system, and a GPAI question for the model

> **The practical rule.**
>
> Two questions decide it: did you change the model, and did you make it available to anyone outside the team that changed it. If both answers are yes, obtain legal advice before shipping. The cost of being wrong is that you owe a documentation and copyright regime you never scoped.

19.5 GPAI controls

BOARD EXTRACT, LIFTABLE

We build on models we did not train, which creates two exposures. First, our providers owe us documentation about their models, and we should be demanding it rather than assuming it exists. Second, we can acquire the obligations of a model provider ourselves simply by fine-tuning a model and making it available beyond the team that built it, with no procurement or board decision involved. We test every modified model against that question before it ships.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
GPA-001	**Record, for every foundation model in use, the provider, the exact model and version, the license, and whether you obtained the Article 53 downstream information pack**	You cannot assess a supply chain you have not enumerated, and the downstream pack is the evidence your provider is meeting its own duties	Engineering with Legal	Foundation model register with per-model license and downstream pack reference	EU: Arts. 3(63), 53 • NIST: MAP • ISO: A.10	ALL / GP / EU	On adoption, then on version change	Model version change, provider change, license change, vendor substituting the underlying model
GPA-002	**Determine and record, per model, whether your own activity makes you a GPAI provider, applying the two-question test of modification and onward availability**	Provider obligations can attach without a procurement decision or a board decision	General Counsel with Engineering	Provider determination memo per modified model, with counsel sign-off where the answer is yes or unclear	EU: Arts. 3(63), 51, 53 • NIST: GOVERN • ISO: 4.3	ALL / GP / EU	Before any fine-tuned model is made available	Fine-tuning, distillation, continued pre-training, internal release beyond the building team, rebranding
GPA-003	**Where you are a GPAI provider, maintain technical documentation covering the training and testing process and evaluation results, and keep it current**	Article 53 requires it to be maintained and produced on request, not assembled when asked	Engineering	Model technical documentation set, versioned	EU: Art. 53 • NIST: MEASURE • ISO: 7.5	P / GP / EU	At release, then on each material change	Retraining, architecture change, evaluation results superseded
GPA-004	**Publish a copyright policy addressing compliance with Union copyright law including reservation of rights, and a public summary of training content on the AI Office template**	Both are standalone obligations that survive the open source carve-out	General Counsel with Engineering	Published copyright policy and training content summary, dated	EU: Art. 53 • NIST: GOVERN • ISO: A.7	P / GP / EU	Before placing on the market	New training corpus, template revision by the AI Office, change of licensing position
GPA-005	**Monitor cumulative training compute against the Article 51(2) threshold and notify the Commission within two weeks of meeting or foreseeing it**	The clock runs from foreseeing the threshold, not from crossing it, and silence is itself a breach	Head of Engineering with Legal	Compute tracking record and, where applicable, the notification with its date	EU: Arts. 51(2), 52 • NIST: GOVERN, MANAGE • ISO: 9.1	P / GP / EU	Continuous during training runs	Training run scaling, delegated act adjusting the threshold, Commission designation

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
GPA-006	**For systemic risk models, perform and document adversarial testing to standardised protocols, and report serious incidents to the AI Office without undue delay**	Article 55 obligations are additional to Article 53 and are enforced by the Commission directly	CISO with Engineering	Evaluation and adversarial testing reports, plus AI Office incident correspondence	EU: Art. 55 • NIST: MEASURE, MANAGE • ISO: A.6	P / GP / EU	Before release, then continuously	New capability, threshold crossing, incident, Code of Practice revision

CHAPTER 20

Technical documentation, conformity assessment and registration

How a high risk system is lawfully placed on the market. This chapter is the assembly instruction for the file, the assessment and the marks.

ENFORCEMENT POINT FOR THIS DOMAIN

The release gate. The technical file, the declaration and the registration entry are preconditions of placing on the market, so the gate is the only place they can be enforced rather than chased.

20.1 The technical documentation

Article 11 requires technical documentation for high risk systems to be drawn up **before** the system is placed on the market or put into service, and kept up to date. The required content is set out in **Annex IV**. It is not a report written at the end; it is an assembly of artifacts your other chapters already produce.

Annex IV expects	Produced in this guide by
A general description of the AI system, its intended purpose, and the persons who developed it	Chapter 6, inventory and intended purpose record
A detailed description of the elements of the system and of its development process	Chapter 6 and Chapter 8, model cards and dataset lineage
Detailed information about monitoring, functioning and control, and performance metrics	Chapter 14, monitoring plan and metric register
The risk management system in accordance with Article 9	Chapter 7, assessments and treatment records
A description of changes made through the lifecycle	Chapter 7 and the change and revalidation policy
A list of harmonized standards applied, in full or in part	Compliance record. Note the verify list on standards publication

Annex IV expects	Produced in this guide by
A copy of the EU declaration of conformity	Section 20.3
A description of the post-market monitoring plan	Chapter 14

> **The reframe that saves months.**
>
> If your program is running, the technical file is largely an index into artifacts that already exist. Organizations that treat it as a document to author, rather than a compilation to assemble, discover at the worst moment that the underlying artifacts were never produced. Build the index early and the gaps become visible while there is still time.

20.2 Conformity assessment

Article 43 sets out the routes. Which applies depends on the category of high risk system and, for some, whether harmonized standards have been applied. Broadly, most Annex III systems follow an **internal control** route in which the provider assesses its own conformity, while certain categories, notably parts of the biometrics area, and the Annex I product route, can require the involvement of a **notified body**.

Confirm your specific route before planning around it. The allocation between internal control and notified body assessment, and the effect of applying harmonized standards, is on the verify list and turns on the precise category your system falls into. Do not assume the lighter route.

A substantial modification to a high risk system after it has been placed on the market can require a fresh conformity assessment. Treat conformity as attached to a configuration, not to a product name.

20.3 Declaration, marking and registration

Step	Provision	What it produces
Draw up the EU declaration of conformity	Art. 47	A signed declaration, retained and available to authorities. It is a statement of responsibility, not a formality

Step	Provision	What it produces
Affix the CE marking	Art. 48	The visible mark, applied only once conformity is assessed and declared
Register in the EU database	Arts. 49, 71	A registration entry before placing on the market or putting into service. Certain deployers also register

The mechanics of the database, including exactly who registers what and at what point, have been operationally fluid and are on the verify list. What is stable is the sequence: assess, declare, mark, register, then place on the market. Doing them out of order is the common failure, and it is visible in the dates on your own artifacts.

20.4 Conformity controls

BOARD EXTRACT, LIFTABLE

Placing a high risk AI system on the European market has a fixed sequence: assess conformity, sign the declaration, apply the CE marking, register, then release. The technical file the law requires is largely an index into artifacts our governance program already produces, so the risk is not writing it but discovering that the underlying artifacts were never created. We build that index at project start rather than before launch.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
CNF-001	**Build the technical file as an index into existing artifacts, with a named owner per Annex IV item and a link to where each lives**	A file authored at the end reveals missing artifacts too late to produce them	Compliance with Engineering	Annex IV index with owner and artifact location per item, showing gaps	EU: Art. 11, Annex IV • NIST: MEASURE • ISO: 7.5	P / HR / EU	At project start, then maintained	New Annex IV item, artifact relocation, system change
CNF-002	**Determine and record the applicable conformity assessment route per system, including whether a notified body is required, and obtain confirmation before planning to the lighter route**	Assuming internal control where a notified body is required loses months and cannot be recovered near a deadline	Compliance with Legal	Route determination per system with the basis recorded	EU: Art. 43 • ISO: 9.2	P / HR / EU	At classification, then on change	Category reassessment, harmonized standard publication, substantial modification
CNF-003	**Draw up and sign the EU declaration of conformity, retain it, and include a copy in the technical file**	Annex IV expects the declaration inside the file, which means the file cannot be complete before conformity is settled	Compliance	Signed declaration, dated, with retention location	EU: Arts. 47, 11 • ISO: 7.5	P / HR / EU	Before placing on the market	Substantial modification, change of intended purpose, new version placed on the market

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
CNF-004	**Sequence assess, declare, mark, register and place on the market, and verify the dates on your own artifacts are in that order**	Out of order execution is visible in your own document dates and is the easiest finding an auditor can make	Compliance	Sequence evidence: dated assessment, declaration, registration and release records	EU: Arts. 43, 47, 48, 49 • ISO: 8.1	P / HR / EU	Per release	New release, re-registration, substantial modification
CNF-005	**Define what constitutes a substantial modification for each high risk system, and require conformity reassessment when the definition is met**	Conformity attaches to a configuration. Without a definition, the trigger never fires	Compliance with Engineering	Substantial modification definition per system, plus reassessment records	EU: Arts. 43, 25 • ISO: 8.1	P / HR / EU	At release, then on each change	Model substitution, purpose change, capability extension, retraining beyond the defined envelope

CHAPTER 21

Jurisdictions beyond the EU, and sector overlays

The EU is the most developed regime, not the only one. This chapter is the map for a portfolio that operates in more than one market, and for regulated industries where sector rules arrive first.

21.1 How to read a multi-jurisdiction position

Three rules prevent the most common errors. **Binding and voluntary are not equivalent** and must not be recorded in the same column. **Sector regulation stacks** on top of horizontal AI law rather than replacing it, and in regulated industries the sector regulator usually reaches you first. And **the map decays**: in the twelve months before publication two of the regimes below changed materially.

21.2 The principal regimes

Jurisdiction	Status	Position as verified August 2026
European Union	Binding	AI Act as amended by the 2026 Digital Omnibus. Annex III high risk from 2 Dec 2027, Annex I from 2 Aug 2028. Article 50 transparency and Article 4 literacy live since 2 Aug 2026. GPAI since 2 Aug 2025
United Kingdom	Binding via regulators	No horizontal statute. Existing regulators apply their own remits. Map your AI to the regulators that already govern you
Council of Europe	Treaty	Framework Convention on AI, Human Rights, Democracy and the Rule of Law. Shapes national law rather than binding firms directly
United States, federal	Executive policy	No comprehensive statute. Preemption has been pushed but not enacted, so state law continues to bind. Do not treat the push as suspending anything
US, California	Binding	Frontier AI transparency, training data disclosure, AI transparency and CPPA automated decision technology rules, phasing across 2026 to 2028

Jurisdiction	Status	Position as verified August 2026
US, Colorado	Binding	IMPORTANT: the 2024 AI Act was repealed and replaced. The successor commences 1 Jan 2027 and is materially narrower, having dropped the duty of care and deployer impact assessments
US, Texas	Binding	In force 1 Jan 2026. Intent based prohibitions, exclusive Attorney General enforcement, and a safe harbour benefit for adopting the NIST AI RMF
US, Illinois and New York	Binding	AI in employment decisions, and independent bias audits for automated employment decision tools in New York City
Canada	No AI statute	IMPORTANT: AIDA lapsed with Bill C-27 in January 2025 and has not been reintroduced. Use privacy law, provincial automated decision rules and financial sector model risk guidance
China	Binding	Algorithm filing and security assessment with the regulator are a precondition of public launch. Content labeling obligations apply to generated material
South Korea	Binding	Framework legislation in force since 22 January 2026, extraterritorial, with a domestic representative requirement above thresholds
Japan	Voluntary	Promotional statute, no penalties. Watch sector regulators layering binding rules on top
India	Mixed	Governance guidelines are voluntary. The IT Amendment Rules on synthetically generated information are binding from 20 February 2026, with labeling and traceability duties. Data protection rules phase through 2027
Singapore	Voluntary	Model AI Governance Framework, including the Agentic edition at v1.5. Influential and the clearest available agentic guidance. No penalty attaches
Brazil and Australia	Proposed	EU style bill and voluntary guardrails respectively. Monitor rather than build to

The default for anywhere without a local AI law.

Govern to the EU AI Act plus ISO/IEC 42001. That combination satisfies substantially every emerging regime and is the safest global baseline, because obligations elsewhere have so far been a subset of, or convergent with, that pair.

21.3 Sector overlays

Where you are regulated, these stack on top of horizontal AI law and usually demand more.

Sector	Instruments that stack	What they add
Financial services	US model risk management supervisory expectations, UK prudential model risk rules, EU operational resilience rules, Canadian model risk guidance, Indian central bank and securities regulator expectations	Model inventory, independent validation, board level model risk appetite, ongoing monitoring, third party resilience registers
Life sciences	Device regulation and predetermined change control plans, electronic records and signatures rules, computerised system validation practice, EU device regulation	Lifecycle validation rather than one time validation, audit trails and record integrity, a pre specified change envelope
Insurance	Model bulletins on AI use, national insurance regulator expectations	A written AI program, third party model oversight, documented testing for unfair discrimination
Employment and HR	Bias audit and notice regimes in several US jurisdictions, automated decision rules	Independent bias audits, candidate notice, human review of adverse automated decisions
Healthcare	Health privacy rules, device pathways where clinical claims are made	Protected health information safeguards in training and inference
Public sector and suppliers	Agency AI use directives	Use case inventories, minimum practices for rights and safety impacting AI, named accountable officers
Critical infrastructure	Network and information security regimes, sector resilience rules	Resilience, incident reporting and supply chain duties stacking on the AI obligations

21.4 Jurisdiction and sector controls

BOARD EXTRACT, LIFTABLE

We operate under more than one AI regime and they do not align. Our position is recorded per market, distinguishes binding law from voluntary frameworks, and is re-verified quarterly because it decays quickly. Where a market has no AI law of its own we govern to the EU AI Act and ISO/IEC 42001 as the baseline. Where we are regulated, sector rules stack on top of AI law rather than replacing it, and they usually reach us first.

ID	Action	Why it matters	Accountable	Evidence artifact	Framework references	Applies when	Frequency or trigger	Validity condition
JUR-001	**Maintain a register of every market in which each system operates or its output lands, resolved against the applicable regimes, and re-verify it quarterly**	Two of nineteen regimes changed materially in six months. A map refreshed annually is wrong for most of the year	AI Gov Lead with Legal	Jurisdiction register with per regime status and the date last verified	EU: Art. 2(1)(c) • NIST: MAP • ISO: 4.1, 4.2	ALL / ALL / ALL	Quarterly	New market, new customer geography, regulatory change, output routing change
JUR-002	**Identify the sector instruments that stack on your AI obligations and map each to the control that already satisfies it**	Sector rules rarely need new controls. They usually need existing controls evidenced to a higher standard	Compliance with the sector SME	Sector overlay map linking each instrument to the controls that discharge it	Sector specific • ISO: 4.1 • NIST: GOVERN	ALL / ALL / ALL	Annually, and on regulatory change	New sector rule, new supervisory expectation, entry into a regulated line of business
JUR-003	**Where a regime is voluntary, record it as voluntary and do not represent adoption as compliance in any internal or external statement**	Recording a voluntary framework as though it were binding misstates your position to your own board	AI Gov Lead	Register showing binding, voluntary or proposed per regime	ISO: 4.2 • NIST: GOVERN	ALL / ALL / ALL	At register creation, then quarterly	Regime changing status, new framework adopted, external claim being drafted
JUR-004	**For any jurisdiction with no local AI law, adopt the EU AI Act plus ISO/IEC 42001 as the governing baseline and record that decision**	An explicit baseline decision prevents each market being negotiated separately	AI Gov Lead with Legal	Baseline decision record per market	EU: whole Act • ISO: 4.1 • NIST: GOVERN	ALL / ALL / ALL	On market entry	New local AI law enacted, material divergence from the baseline identified

APPENDIX A

Glossary

Term	Definition
Action space	The set of actions an agent is technically able to perform. With autonomy, one of the two axes determining agent risk
AIBOM	AI bill of materials. Enumeration of models, datasets, libraries, frameworks and tool servers with versions
Annex I route	High risk classification via Article 6(1), for AI as a safety component of, or as, a product under listed harmonization legislation
Annex III route	High risk classification via Article 6(2), for standalone systems in one of eight enumerated use areas
Attestation	A dated statement by a named person that a control operated, valid until a named voiding event
Completeness	The property that no records are missing. Distinct from tamper evidence and not addressed by it
Deployer	Whoever uses an AI system under their own authority in a professional capacity
Emission	The writing of a record by the component that authorizes the action
Enforcement point	A component that must be traversed for an action to occur
Evidenced	A control state meaning the artifact exists, is current, is owned, and is producible within one working day
Fail closed	Emission failure blocks the action, rather than the action proceeding unrecorded
GPAI	General purpose AI model. Parallel obligation track under Articles 51 to 55
Observer	A component that watches actions occur without being required for them
Provider	Whoever develops an AI system, or has it developed, and places it on the market under their own name or trademark

Term	Definition
Role migration	A deployer becoming a provider through rebranding, substantial modification, or change of intended purpose
Validity condition	The enumerated change events that void an attestation, with the watcher for each
Watcher	The mechanism that detects a voiding event. A condition without one is an intention
Write path	The design decisions determining which code paths emit, synchronicity, failure behavior, sequencing and overflow. Sets the completeness ceiling

APPENDIX B

Source register

Instrument	Status	Official source
Regulation (EU) 2024/1689 (AI Act)	Binding law	eur-lex.europa.eu/eli/reg/2024/1689/oj
AI Act Service Desk, Article 2	Commission	ai-act-service-desk.ec.europa.eu/en/ai-act/article-2
AI Act Service Desk, Article 5 FAQ	Commission	ai-act-service-desk.ec.europa.eu
AI Act Service Desk, Article 73	Commission	ai-act-service-desk.ec.europa.eu/en/ai-act/article-73
Commission draft guidance on reporting serious incidents, 26 Sep 2025	Draft guidance	Commission AI Act pages
Commission Guidelines on prohibited AI practices, C(2025) 884 final, 4 Feb 2025	Commission guidance	Commission AI Act pages
Draft Commission guidelines on high risk classification, 19 May 2026	Draft guidance	Commission AI Act pages
2026 Digital Omnibus amending instrument	Binding, amending	EUR-Lex, cite the amending act
NIST AI 100-1, AI RMF 1.0	Voluntary framework	nist.gov/itl/ai-risk-management-framework
NIST AI 600-1, Generative AI Profile	Voluntary profile	NIST publications
NIST AI RMF Playbook	Companion resource	airc.nist.gov/airmf-resources/playbook
ISO/IEC 42001:2023	Certifiable standard	iso.org
ISO/IEC 42005, 42006, 23894, 5338, 38507, 12792, 22989	Standards	iso.org
Singapore Model AI Governance Framework, GenAI edition	Voluntary	IMDA and AI Verify Foundation

Instrument	Status	Official source
Singapore MGF for Agentic AI, v1.5, 20 May 2026	Voluntary	imda.gov.sg
AI Verify	Testing framework and toolkit, not a governance framework	AI Verify Foundation
OWASP Top 10 for LLM Applications	Community reference	owasp.org
MITRE ATLAS	Community reference	atlas.mitre.org

APPENDIX C

Verify before publication

Confirm each item against a primary source before relying on it operationally.

Resolved during preparation, no further action

Article 2(1) actor categories and territorial hooks. Article 5(1)(a) to (h) prohibited practices, and the workplace and education scope of Article 5(1)(f) including the customer versus employee distinction. Articles 6(1) to 6(4) classification, the profiling bar and the pre-market documentation requirement. Articles 9(1) and 9(2). Articles 10(1) to 10(6). Article 73(1) to (4) reporting deadlines of 15, 10 and 2 days with their triggers, and the Article 3(49) definition. Article 99(3), (4), (5) and (6) penalty tiers including the SME reversal. NIST AI RMF 1.0 four-function structure. Singapore MGF Agentic version 1.5 dating.

EU AI Act, provisions still open

Article 9 paragraph numbering for the testing provisions and the vulnerable groups provision. Article 12 wording, specifically whether it obliges logging capability or the act of logging. Articles 19 and 26 log retention durations. Articles 49 and 71 registration mechanics. Article 50 paragraph allocation and exemption wording. Article 72 Commission template adoption. Article 53 GPAI copyright policy obligation. Whether Article 5(1)(c) social scoring reaches private-sector actors as well as public authorities.

2026 amendments and guidance

Final status and content of the Digital Omnibus, including the Annex III date of 2 December 2027 and the Annex I date of 2 August 2028. Reported extension of Article 5 to further practices from 2 December 2026. Whether the 19 May 2026 draft high risk classification guidelines are final. Whether the 26 September 2025 draft guidance on serious incident reporting is final, including the simplified route where equivalent sector reporting exists. Whether harmonized standards underpinning high risk conformity have been published.

NIST AI Risk Management Framework

Function-level references (GOVERN, MAP, MEASURE, MANAGE) are verified against NIST AI 100-1 and are used throughout. Subcategory identifiers were removed except where individually confirmed. If subcategory-level citation is required, every identifier must be checked against Tables 1 to 4 of NIST AI 100-1.

ISO/IEC standards

Every clause and Annex A group reference, against licensed copies. No normative text is reproduced anywhere in this document, and the Chapter 18 readiness checklist is written against the audit process rather than the standard's structure for that reason.

Singapore Model AI Governance Framework

MGF GenAI dimension names as used in the Chapter 5 crosswalk, which are indicative mappings rather than IMDA statements. Agentic v1.5 dimension wording. AI Verify current scope and version.

APPENDIX D

Change log

Version	Date	Change
1.0	August 2026	Initial release. Regulatory positions verified August 2026 against primary and official Commission sources. NIST references made at function level following verification of AI RMF 1.0 structure.

APPENDIX E

Stay current

This document is versioned rather than permanent, because the positions in it change. The current release, and the companion 27 sheet workbook, are always at akhawat.com/toolkits.

Subscribe to the CXO Intelligence Series	linkedin.com/newsletters/7475953227239419904
All articles in the series	akhawat.com/writing
Follow the author on LinkedIn	linkedin.com/in/prashantakhawat
Toolkit and current release of this guide	akhawat.com/toolkits

The companion article, *Three Frameworks, One System: The Convergence Playbook for AI Governance*, sets out the argument this guide operationalises and is available with this guide at akhawat.com/toolkits.

APPENDIX F

Disclaimer

This document is an educational and operational reference. It is not legal advice and does not create a lawyer-client relationship. It states positions current as of the date on the cover, in a field that changes materially over periods of months. In the twelve months preceding publication, Colorado repealed and replaced its 2024 AI Act and Canada's Artificial Intelligence and Data Act lapsed without reintroduction.

References to ISO/IEC standards are by number, clause and intent only. No normative or copyrighted standard text is reproduced. Licensed copies are required to implement against any of them. The Chapter 18 readiness checklist is written against the audit process rather than the structure of the standard, for that reason.

Verify every position against the primary source for your jurisdiction, and obtain professional advice before making compliance decisions. The author accepts no liability for reliance on this document.

Copyright 2026 Prashant Akhawat. Free to use and adapt inside your own organization. Not for resale, redistribution or white-labeling without written permission. Current release at akhawat.com/toolkits