



COADF

Compliance-Oriented AI Development Framework

COADF is publicly documented as a development framework. No authority has assessed, audited or endorsed it.

Version 2.2

September 2026

COADF version 2.2, edition r2

COADF · Compliance-Oriented AI Development Framework

Version 2.2 · September 2026 · COADF version 2.2, edition r2

This edition

Framework version: 2.2

Edition revision: r2

Language: en

Publication date: 2026-09-07

Public source commit: ed15f765e68e44aa9e4a83230b8e93e335b08308

Public projection hash: d74dc0091e6eac24f5fe83f3307b9f3c5b89e0f497dca4722444d2f172b22e2b

COADF is publicly documented as a development framework. No authority has assessed, audited or endorsed it.

Proposed licence: CC BY-SA 4.0 for the text, MIT for the schemas. Not yet granted.

IP and publication status

A Brazilian patent application relating to certain implementation aspects is pending before the Brazilian Patent and Trademark Office (INPI), application BR 10 2026 020627 0. No patent has been granted. The public COADF materials describe architectural principles and practices. They do not represent the legal scope of the pending application. Certain implementation details remain outside the public edition.

How to cite

Proposed attribution: COADF, by L. C. Hogrefe (AnyLAI), version 2.2, 2026.

COADF, by L. C. Hogrefe (AnyLAI), version 2.2, 2026

Overview



COADF is a development method for AI systems that have to hold up under European regulation.

It states eight architecture principles, each tied to the obligation it was written against, with governance templates and a machine-readable model for reporting which controls a project actually exercises.

It is the method behind AnyDPP, and it is applied to the project that produced it, which is where the report on this site comes from.

Where COADF sits

COADF is the AI development core of the Semantic Trust Framework. The runtime path a decision travels is the Trust Corridor. The surface where a person sees and resolves what the system stopped is the Trust Console. COADF is the part that says how the software is built.

The eight principles

Some principles are published in full. Others state the principle and stop, because the operative detail belongs to an application that has been filed and not published. Every section says which of the two it is, rather than leaving a reader to notice the difference.

-
- **P-1 · Deterministic first, probabilistic quarantined** · Published in full. Deterministic processing by default. A probabilistic component is bounded and never reaches an output directly.
 - **P-2 · Confidence-gated output** · Published as a principle. Confidence attaches to one attribute at a time, and low confidence never reaches a published output.
 - **P-3 · Human review by architecture** · Published as a principle. Where confidence is insufficient a person decides, at attribute level, against the source.
 - **P-4 · One trace, end to end** · Published in full. One identifier links ingestion, every step and the published output, and the chain exports.
 - **P-5 · Disclosure of machine extraction** · Published in full. An output carrying machine-extracted data says so, and names which attributes.
 - **P-6 · The fence system** · Published as a principle. Guardrails are automated checks that block a release, not advice in a document.
 - **P-7 · Standards dependency isolation** · Published in full. An external service can raise confidence. None of them can gate an output.
 - **P-8 · Policy as data, with graduated autonomy** · Published as a principle. Rules are data, the engine does not change when a rule changes, and a decision is bound to the rule version that produced it.
-

Two things COADF is not

It is not a certification. Nothing here is issued, awarded or withdrawn by anybody, and running it produces no mark of any kind.

It is not a norm. It does not stand against ISO/IEC 42001, the NIST AI Risk Management Framework or an audit assurance architecture. It describes how to build the evidence those ask for.

Licence, proposed and not yet granted

The proposal is Creative Commons Attribution ShareAlike 4.0 International for the document text and this site copy, so that a document adapted from it stays open, and the MIT licence for the published schemas and the artifacts generated from them.

MIT rather than Apache 2.0, deliberately: Apache section 3 grants patent rights expressly, and there is a filed application in this family. That is a reason to prefer MIT here; it is not a reason to think MIT settles any patent question, because it does not.

The proposed CC BY-SA licence covers the text only. It grants no patent rights and does not impose share-alike obligations on independent implementations.

Proposed attribution: COADF, by L. C. Hogrefe (AnyLAI), version 2.2, 2026.

The grant the framework document offers enumerates P-1 through P-7 and says nothing at all about P-8. That silence is a fact about the text, not a decision: it is not stated as an exclusion and it is not stated as an inclusion, and the scope over P-8 is unresolved. It is written here rather than left for a reader to notice by comparing lists.

No grant is made by this page. The notice above is the proposed wording, held for the founder's decision on scope, and the file-by-file inventory it applies to is prepared alongside it. A licence, once granted publicly, is not taken back by calling it provisional, which is precisely why it is not granted here first and decided afterwards.

Principles



One section per principle. Each names the obligation it was written against. An anchor here states alignment and nothing more: it is the article the principle was designed to answer to, never a statement that this project is subject to it.

P-1 .



Deterministic first, probabilistic quarantined

· Published in full

Everything that reaches a person, a document or another system is deterministic by default: parseable, reproducible, and reconstructable from its inputs. Machine learning components are held to bounded tasks. The boundary is explicit in the code rather than in a diagram: a probabilistic component returns a value, a confidence and the method the value was extracted by, never a bare value. A bare value is the failure this principle exists to prevent, because a bare value cannot be gated, disclosed or reviewed.

Alignment:

Regulation (EU) 2024/1689, Article 9 on the risk management system and Article 15 on accuracy and robustness.

Source: Regulation (EU) 2024/1689, Article 9 (risk management system) and Article 15 (accuracy, robustness and cybersecurity) ([primary source](#)) · checked on 7 September 2026

P-2 .



Confidence-gated output

· Published as a principle

Confidence attaches to a single attribute, not to a document. A document is not right or wrong as a whole; each fact it carries is separately well or poorly evidenced, and one score for the file hides exactly that. Three levels exist. The lowest never reaches a published output; it goes to a person instead. Where there is no evidence at all there is no value: the output says the value is not available, and never estimates it.

Not published:

The rules that map an extraction method to a level, and the arithmetic that combines several sources into one, are not published here. They belong to a filed application. A framework loses nothing by describing what it stops and staying silent on the arithmetic that stops it.

Alignment:

Regulation (EU) 2024/1689, Article 14 on human oversight and Article 15 on accuracy.

Source: Regulation (EU) 2024/1689, Article 14 (human oversight) and Article 15 (accuracy, robustness and cybersecurity) ([primary source](#)) · checked on 7 September 2026




Figure · The confidence gate.

Three lanes of items move toward a gate. One passes through. One passes carrying a flag. One stops at the gate and moves across to a figure standing beside it. A fourth lane is empty, and nothing travels along it.

Three lanes reach the gate and a fourth is empty. One item passes. One passes flagged. One stops and goes to a person. The empty lane is the attribute with no evidence: it produces nothing at all, rather than a guess. The picture shows the shape and names no level and no number, which is the boundary this framework publishes at.

P-3 ·



Human review by architecture

· Published as a principle

Review is triggered by confidence, not by a schedule and not by somebody noticing. It happens at attribute level: the reviewer sees one fact and the passage of the source it came from, side by side, and decides about that fact. A rejected value leaves the attribute empty. The system never fills it with a second guess, because the reason for the review was that the evidence was insufficient, and a guess does not repair that.

Not published:

How the review queue is built and ordered, and what is demanded of a reviewer at the moment a decision is written, are not published here.

Alignment:

Regulation (EU) 2024/1689, Article 14 on human oversight.

Source: Regulation (EU) 2024/1689, Article 14 (human oversight) ([primary source](#)) · checked on 7 September 2026

P-4 .



One trace, end to end

• Published in full

A single identifier is created when a document is taken in, and it travels with everything that happens afterwards: every processing step, every confidence assessment, every human decision, and the published output. The trail is append-only, so a later entry corrects an earlier one and nothing is rewritten. One query reconstructs the chain from a published output back to the documents it rests on, and the chain exports as JSON for somebody who did not build the system.

Alignment:

Regulation (EU) 2024/1689, Article 12 on record-keeping and Article 13 on transparency towards deployers. For product passports, Regulation (EU) 2024/1781, Article 13 requires the Commission to set up a registry storing at least the unique identifiers.

Source: Regulation (EU) 2024/1689, Article 12 (record-keeping) and Article 13 (transparency and provision of information to deployers) ([primary source](#)) · checked on 7 September 2026

Source: Regulation (EU) 2024/1781, Article 13(1) ([primary source](#)) · checked on 7 September 2026




Figure · The trace.

A single thread enters at the left and stays visible as it passes through four stages to a published document at the right. A query then travels back along the same thread to where it started.

One thread enters at intake and is still the same thread at the published output. The return journey is the point: any published fact can be walked back to the documents it rests on, in one query.

P-5 .



Disclosure of machine extraction

· Published in full

Every attribute carries the method it was extracted by. Where any attribute in a published output was read by optical character recognition or by a language model, the output carries a visible line saying so and naming which attributes it applies to. The line sits on the document a reader is looking at, not in metadata a reader would have to go and find.

Alignment:

Regulation (EU) 2024/1689, Article 50 on transparency obligations for providers and deployers of certain AI systems.

Source: Regulation (EU) 2024/1689, Article 50 (transparency obligations for providers and deployers of certain AI systems) ([primary source](#)) · checked on 7 September 2026

P-6 .



The fence system

· Published as a principle

A guardrail that lives in a document is advice. A guardrail that runs in the pipeline and blocks a release is a fence. Fences fall in four categories: data, architecture, copy and process. Each has an identifier, a rule, and an enforcement method. A fence earns its place only by a proof of teeth: the real defect is planted, the fence fires, and the plant is removed byte for byte. A fence nobody has watched fail is a comment with a test runner attached.

Not published:

The vocabulary the publication fence matches on is not itself published. A list of what is being watched is a map of what is being protected, and printing the map defeats the fence.

Alignment:

Regulation (EU) 2024/1689, Article 9 on the risk management system and Article 17 on the quality management system.

Source: Regulation (EU) 2024/1689, Article 9 (risk management system) and Article 17 (quality management system) ([primary source](#)) · checked on 7 September 2026




Figure · The fence.

A change moves toward a deployment target and strikes a wall, which lights up with the mark of the rule that stopped it. A second change moves along the same path and passes through.

A fence blocks a release and shows which rule did it. The second change passes, which is the half that matters: a fence that stops everything is a wall, and a fence nobody has watched stop something is a comment.

P-7 .



Standards dependency isolation

· Published in full

Classification systems, terminology databases and external validation services live in adapter modules, marked as such, outside the core data model. An external service can raise the confidence of an attribute. None of them can gate an output: when a service is unavailable the system produces less confident evidence rather than nothing at all, and when a licence changes the core is untouched. Licensed content is marked in its metadata and kept apart from the open core.

Alignment:

Regulation (EU) 2024/1689, Article 9 on the risk management system.

Source: Regulation (EU) 2024/1689, Article 9 (risk management system) ([primary source](#)) · checked on 7 September 2026

P-8



Policy as data, with graduated autonomy

· Published as a principle

The rules a decision is measured against are data. The engine that evaluates them does not change when a rule changes. A decision record is bound to the version of the rule set that produced it, which extends the trace of P-4 into rule provenance. Autonomy is a property of the rules rather than of the code: how much the system may do without a person is set per jurisdiction and per risk class, and it never exceeds the ceiling the law sets there. Raising it where the law allows is a rule change with a recorded legal basis, never a fork of the code.

Prior art:

Signed, reloadable rule sets with version-linked decision records already exist, and Open Policy Agent is the obvious example. Nothing here presents that idea as new. The contribution here is the practice of binding it into the rest of this architecture.

Not published:

How a rule set is sealed, how the seal is checked, how the ceiling is enforced, and how any of that meets the confidence gate, are not published here.

Alignment:

Regulation (EU) 2024/1689, Article 12 on record-keeping and Article 17 on the quality management system. Where Article 14 applies, effective human intervention stays wired in whatever the rules say.

Source: Regulation (EU) 2024/1689, Article 12 (record-keeping), Article 17 (quality management system) and Article 14 (human oversight) ([primary source](#)) · checked on 7 September 2026



Figure · Policy as data.

A sealed block arrives at an engine block that does not change shape. A record leaves the engine carrying a small stamp that matches the seal.

Rules arrive as data and the engine is not rebuilt to receive them. The record that leaves carries the version stamp of the rules that produced it. The interior of the seal is not shown, and is not published.

Honest absence

Inside P-2, and stated separately because it is the rule most often broken: no evidence produces no value. Not a default, not a median, not a plausible figure. An empty attribute is a true statement about what is known, and a filled one that nobody can trace is not.

Candidate principles

Four practices that behave like principles and are not yet numbered as such. They are published because they are in use, and a practice in use that nobody wrote down is a practice that decays.

Claim discipline

Every regulatory sentence resolves to a primary source and the date somebody read it. Every sentence about our own behaviour resolves to deployed behaviour or to a named test. When a sentence turns out to have neither, there are two honest moves: build the thing, or soften the sentence. Citing a source nobody opened is not among them. This has no value to an application and a great deal of value to a reader.

Non-override

An authenticated adverse fact from an official source cannot be attested away. A finding the system derived itself can be resolved by a named person, and the resolution is recorded with who made it and what they saw. The distinction is the entire point: a person may overrule an inference of ours and may not overrule a register.

Evidence semantics

One fact carries several readings at once and they never collapse into a single number. Is the source genuine. What does it actually cover. What did it return. How recent is it. Where did it come from. Was a person involved. Which duty does it speak to. A high reading on one and a low reading on another is not a contradiction to be averaged away; it is the information.

Not published:

The states each reading moves through, the rules that move them, and the way a narrow coverage limits the rest, are not published here.




Figure · The seven readings.

One fact at the centre opens into seven separate dials arranged around it. The dials settle at different levels and never merge into one.

One fact, seven readings, kept apart.

Disclosure control

A fence keeps inventive matter that has not been filed off every public surface, including built frontend bundles, and a register records what is being held and the condition under which it stops being held. That the discipline exists is worth stating. What it guards is not.

Communication red lines

Binding on every framework and product text. Each row is a sentence that is easy to write and impossible to support.

Say: Audit-ready evidence that management met its oversight duty.

Never: Protection from personal liability.

Why: Oversight duties are statutory and are not waived by a tool.

Say: Human oversight by design, human on the loop.

Never: Human out of the loop.

Why: Where Article 14 of Regulation (EU) 2024/1689 applies, effective human intervention is the requirement.

Source: Regulation (EU) 2024/1689, Article 14 ([primary source](#)) · checked on 7 September 2026

Say: Verifiably reduces the risk of a fabricated value.

Never: Eliminates hallucination.

Why: An architecture lowers a risk. It does not remove a property of the model class.

Say: A regulation-wired framework with a runtime enforcement path.

Never: The first AI-native enterprise architecture framework.

Why: There is prior art, and a first-of-its-kind claim invites the one argument nobody needs.

Governance templates



Three templates and one description. All four are practice rather than mechanism, and all four are published in full.

AI agent charter

When agents write production code they need the governance a developer has, and a little more. The ladder is about how much an agent may do before a person looks. The list under it is about what nobody delegates, at any rung.

The autonomy ladder

1. **L1 · Supervised.** The agent proposes and a person reads every line before it merges. The default for every new agent.
2. **L2 · Semi-autonomous.** The agent implements inside a defined ticket scope and a person reviews the pull request. Reached after three L1 cycles with no fence violation.
3. **L3 · Autonomous within guardrails.** The agent implements, tests and opens the pull request; a person reads the test results and the guard output. Requires mature fence coverage and a track record. Rare, and it should stay rare.

Never autonomous

A person decides these, at every rung of the ladder.

- Database schema changes.
- Security configuration: authentication, secrets, network.
- External interface contracts, where the change breaks a caller.
- Changes to an output format that a regulation reads.
- Deletion of anything in the audit trail.
- Changes to the fence system itself.

The agent's source of truth

Architecture documentation first, then the key registry, then the fence system, then the ticket scope. Where two of them disagree, the higher one wins, and the disagreement is written down rather than resolved silently.

Audit trail schema

The record written at every step, published in full. A trail that cannot be read by somebody who did not build the system is not a trail, so the field names are part of the framework rather than an implementation detail.

Field	Meaning
trace_id	The identifier of P-4, the same value on every entry that belongs to one transaction.
timestamp	When the step happened, as a timestamp with its offset.

Field	Meaning
event_type	Which kind of step: intake, extraction, validation, confidence assessment, human verification, publication, access.
actor	Who acted, and of what kind: the system, a person, or an agent.
input	What the step was about: the document, the attribute.
output	What the step produced: the value, its confidence, the extraction method.
decision	Where a person acted: what they did and the reason they gave.
source_hash	A digest of the source the step read, so the same input can be recognised later.
immutable	The entry is never edited. A correction is a new entry that supersedes it.

Fence registry, a published subset

A published subset of the project's own registry, with the identifiers the project actually uses. Each entry carries an identifier, the rule, how the rule is enforced and its current status. The full registry is longer and stays internal: some of its entries name what they protect, and a list of what is watched is a map of what is guarded.

The identifiers are the operational ones. An earlier edition of the framework document printed an eight-row illustrative table under the same numbers with different rules; that template is retired, every one of its rules was carried across to the identifier that now holds it, and nothing was lost in the renumbering.

Fence	Rule	Enforcement	Status
F-03	Values derived by a language model always carry the method they were extracted by and never reach a published output without human verification.	Automated check, on every change	■ Enforced
F-04	A required attribute with no evidence is never fabricated and never defaulted.	Automated check, on every change	■ Enforced
F-05	No real personal or company documents in code, fixtures, tests or commits. Scope: The guard covers export-document access keys and issuer data. It does not cover every identifier shape, and the residue is carried internally as its own gap.	Automated check, on every change	■ Enforced
F-08	A passport carrying a low-confidence or missing required attribute is never published; the publication gate is code, not a habit.	Automated check, on every change	■ Enforced
F-09	No database schema change without a migration and a working downgrade.	Automated job, on every change	■ Enforced

Fence	Rule	Enforcement	Status
F-13	The tenant and trace identifiers travel with every request and every stored row. Scope: The request-path half is tested. The stored-row half is not asserted column by column, and is carried internally as its own gap.	Automated check, on every change	■ Enforced
F-14	Code and identifiers in English; reader-facing copy in five locales, with every key present in every one.	Automated check, on every change	■ Enforced
F-16	Interface copy never asserts legal conformity or an unverified fact; an attestation renders as an attestation and never as a verification.	Automated check, on every change	■ Enforced
F-18	A trace identifier on every console interface response, errors included; an incoming one is honoured rather than reminted.	Automated check, on every change	■ Enforced
F-22	Every factual claim on a public surface resolves to a primary source or to a named test.	Automated check, on every change	■ Enforced
F-23	The running site loads only same-origin resources and sets exactly one cookie, for language.	Automated check, on every change	■ Enforced
F-37	Every crawled public page meets WCAG 2.2 AA in a real browser.	Automated job, on every change	■ Enforced
F-02	No new infrastructure service beyond the agreed set, and no new dependency without a named section in the change request.	Named person, at merge	□ Manual
F-10	No merge with failing checks.	Named person, at merge	□ Manual
P-4-not-null-columns	Every stored row carries the tenant and trace identifiers as required columns.	Nothing today	□ Gap
P-5-disclosure-rendered	A published output whose attributes include a machine-extracted one always renders the disclosure line.	Nothing today	□ Gap

The confidence engine, in words

The engine is given an attribute, a value, the method the value was extracted by, and the documents it came from. It returns the attribute, the value, a confidence level, the provenance and whether a person verified it. Three levels exist, the lowest never reaches a published output, and where evidence is absent the engine returns nothing at all rather than a level.

Not published:

The table that maps an extraction method to a level, and the rules that combine several sources, stay in the internal edition. What is above is the description a reader needs in order to understand the architecture. It is deliberately not a specification anybody could build from.

Conformance



A conformance report says which controls a project exercised and what their evidence returned. It is self-attested by whoever runs it. It is not an audit, it produces no mark, and it is worth exactly as much as the evidence behind it.

How a report is produced

Four things make a row, and a row without all four is not reported as enforced.

1. **The control.** One sentence about behaviour, tied to the principle it belongs to. Not a policy and not an intention: something that is either true of the running system or not.
2. **The evidence.** What decides the row. For an enforced control that is an automated check; for a manual one it is a named person performing a written procedure.
3. **The run.** One identified execution. A row is enforced only on the strength of what happened in that run, never on the strength of what usually happens.
4. **The status.** Enforced, manual or gap. Nothing else, and no shading between them.

The three statuses

- **Enforced.** An automated check carries the rule, and it ran in the identified run and passed.
- **Manual.** The control holds through a procedure a named person performs. Real, not automated, and reported as what it is.
- **Gap.** Nothing enforces it today. Named, with what closing it would take, rather than omitted.

The anti-theatre rule

A control is never reported as enforced unless its evidence ran in the identified run and passed. A previous success, a skipped check or a selection that ran no tests at all is not current verification. A report that calls a control enforced because somebody remembers it passing is a report about memory, and running a report successfully is not a certificate that anything is implemented.

What this report is, and is not

It is a self-attested internal assessment with a limited public evidence summary. It cannot be reproduced from this page: the evidence runs against a repository that is not public, so no command here is offered as one a reader could run. It is not an audit, not a certification and not a conformity assessment, and no authority has assessed, audited or endorsed it.

This project's own assessment

The rows below are a reviewed public projection of an internal assessment of this repository. They are a SUBSET, chosen because each row can be stated publicly without naming what a guard protects. The internal assessment is larger and is not published.

Which assessment this is

Framework version: 2.2

Assessed on: 7 September 2026

Run: 34152270255

Produced by: the COADF system work, not the site work. The site did not re-execute the evidence; it projects it, and names whose it is.

The same set of controls, run on a workstation rather than in continuous integration, was red for two controls that need a database the workstation does not have. Both executed and passed in the run named above. Both runs are recorded internally; neither is hidden to improve the look of this page.

Disclosed scope:

16 of 57 controls in the internal assessment. Totals on this page are over the disclosed scope only and are not the totals of the internal assessment. The controls left out are left out because their statements name protected subject matter, not because of their result.

12 enforced, 2 manual, 2 gap, over 16 disclosed controls.

Control	Principle	What it requires	How it is enforced	Status
F-03	P-1	Values derived by a language model always carry the method they were extracted by and never reach a published output without human verification.	Automated check, on every change	■ Enforced
F-04	P-2	A required attribute with no evidence is never fabricated and never defaulted.	Automated check, on every change	■ Enforced
F-05	P-6	No real personal or company documents in code, fixtures, tests or commits. Scope: The guard covers export-document access keys and issuer data. It does not cover every identifier shape, and the residue is carried internally as its own gap.	Automated check, on every change	■ Enforced

Control	Principle	What it requires	How it is enforced	Status
F-08	P-2	A passport carrying a low-confidence or missing required attribute is never published; the publication gate is code, not a habit.	Automated check, on every change	■ Enforced
F-09	P-6	No database schema change without a migration and a working downgrade.	Automated job, on every change	■ Enforced
F-13	P-4	The tenant and trace identifiers travel with every request and every stored row. Scope: The request-path half is tested. The stored-row half is not asserted column by column, and is carried internally as its own gap.	Automated check, on every change	■ Enforced
F-14	P-6	Code and identifiers in English; reader-facing copy in five locales, with every key present in every one.	Automated check, on every change	■ Enforced
F-16	P-6	Interface copy never asserts legal conformity or an unverified fact; an attestation renders as an attestation and never as a verification.	Automated check, on every change	■ Enforced
F-18	P-4	A trace identifier on every console interface response, errors included; an incoming one is honoured rather than reminted.	Automated check, on every change	■ Enforced
F-22	Candidate · Claim discipline	Every factual claim on a public surface resolves to a primary source or to a named test.	Automated check, on every change	■ Enforced
F-23	P-6	The running site loads only same-origin resources and sets exactly one cookie, for language.	Automated check, on every change	■ Enforced
F-37	P-6	Every crawled public page meets WCAG 2.2 AA in a real browser.	Automated job, on every change	■ Enforced
F-02	P-6	No new infrastructure service beyond the agreed set, and no new dependency without a named section in the change request.	Named person, at merge	□ Manual
F-10	P-6	No merge with failing checks.	Named person, at merge	□ Manual

Control	Principle	What it requires	How it is enforced	Status
P-4-not-null-columns	P-4	Every stored row carries the tenant and trace identifiers as required columns. What closing it would take: A test over the data model rather than over a running database, asserting the requirement column by column.	Nothing today	☐ Gap
P-5-disclosure-rendered	P-5	A published output whose attributes include a machine-extracted one always renders the disclosure line. What closing it would take: A render test over an output carrying a machine-extracted attribute, with its own proof of teeth. The rendering exists; nothing fails when it is removed, so the control is a habit rather than a rail.	Nothing today	☐ Gap

Disclaimer

COADF is a publicly documented development framework. It is not a certification, not an audit standard and not a conformity assessment scheme. No authority has assessed, audited or endorsed it. A conformance report is self-attested by whoever runs it, and describes which controls were exercised and what their evidence returned. The name describes the orientation of the development method toward regulatory requirements; it never asserts that any system, output or shipment is compliant.

Regulatory map



The European instruments a system built this way is likely to meet. Every row was read at its primary source on the date the row states, and every row links to that source. This is a map and not advice: nothing here says which of these applies to any particular system.

What this version corrects

The July 2026 edition of this map carried four rows that its own sources do not support. They are corrected here rather than quietly edited, and each correction is a link to the text that settles it.

The registry row put one date on three different events. Article 13(1) of Regulation (EU) 2024/1781 set 19 July 2026 as the deadline by which the Commission was to set the registry up. The Commission announced on 20 July 2026 that the registry was live, which is an announcement of operational availability and not a date in either instrument. Implementing Regulation (EU) 2026/1778 of 16 July 2026 lays down the implementation arrangements; it was published on 17 July 2026 and enters into force on the twentieth day after that. Three dates, three different things, and the July edition described one as the other.

Source: Regulation (EU) 2024/1781, Article 13(1), read with Article 24 of Implementing Regulation (EU) 2026/1778 ([primary source](#)) · checked on 7 September 2026

The due diligence row was written against the original directive alone. Directive (EU) 2026/470 of 24 February 2026 moved the application of those measures to 26 July 2029, set 26 July 2028 as the transposition deadline for its due diligence article, raised the scope thresholds on the principal route for EU companies to more than 5 000 employees on average and a net worldwide turnover above EUR 1 500 000 000, other categories keeping their own criteria, and left the measures for Article 16 to apply for financial years starting on or after 1 January 2030. A separate point, because it is the one most often got wrong: the amending directive removes harmonised Union conditions of civil liability, which is not the same as removing civil liability. National law continues to apply.

Source: Directive (EU) 2026/470, Article 4(2) and Article 5 ([primary source](#)) · checked on 7 September 2026

The critical raw materials row stated a flat start date. Article 29(1) of Regulation (EU) 2024/1252 sets the later of 24 May 2027 and two years after the entry into force of the calculation and verification act under Article 29(2). That act was due by 24 May 2026 and is not published, so the start date is not yet fixed.

Source: Regulation (EU) 2024/1252, Article 29(1) and (2) ([primary source](#)) · checked on 7 September 2026

The row on a scheme for passport service providers described a Commission preparation and a year. Recital 40 of Regulation (EU) 2024/1781 says the Commission could carry out an impact assessment to investigate whether such a scheme is appropriate. That is the whole of what the text supports.

Source: Regulation (EU) 2024/1781, Recital 40 ([primary source](#)) · checked on 7 September 2026

In force and phasing in

EU AI Act

Regulation (EU) 2024/1689

Applies from 2 August 2026. Chapters I and II applied from 2 February 2025; Chapter III Section 4, Chapters V, VII and XII and Article 78 from 2 August 2025; Article 6(1) and the obligations that go with it from 2 August 2027.

Source: Regulation (EU) 2024/1689, Article 113 ([primary source](#)) · checked on 7 September 2026

Cyber Resilience Act

Regulation (EU) 2024/2847

Applies from 11 December 2027. Article 14, the reporting duty, applies from 11 September 2026, and Chapter IV from 11 June 2026.

Source: Regulation (EU) 2024/2847, Article 71 ([primary source](#)) · checked on 7 September 2026

Ecodesign for Sustainable Products Regulation

Regulation (EU) 2024/1781

Framework in force. Article 13(1) required the Commission to set up the digital product passport registry by 19 July 2026; Implementing Regulation (EU) 2026/1778 of 16 July 2026 lays down its implementation arrangements. Delegated acts per product category follow.

Source: Regulation (EU) 2024/1781, Article 13(1), read with Implementing Regulation (EU) 2026/1778 ([primary source](#)) · checked on 7 September 2026

Deforestation Regulation

Regulation (EU) 2023/1115

Obligations bind medium and large operators from 30 December 2026 and micro and small operators from 30 June 2027, after the amendment by Regulation (EU) 2025/2650.

Source: Regulation (EU) 2023/1115, Article 38(2) and (3), as replaced by Regulation (EU) 2025/2650 ([primary source](#)) · checked on 7 September 2026

Battery Regulation

Regulation (EU) 2023/1542

Applies from 18 February 2024. Article 11 applies from 18 February 2027; Article 17 and Chapter VI from 18 August 2024; Chapter VIII from 18 August 2025.

Source: Regulation (EU) 2023/1542, Article 96 ([primary source](#)) · checked on 7 September 2026

General Data Protection Regulation

Regulation (EU) 2016/679

Applies since 25 May 2018.

Source: Regulation (EU) 2016/679, Article 99(2) ([primary source](#)) · checked on 7 September 2026

NIS2 Directive

Directive (EU) 2022/2555

Member States were to adopt the measures by 17 October 2024 and to apply them from 18 October 2024. National transposition is a Member State matter and varies.

Source: Directive (EU) 2022/2555, Article 41(1) ([primary source](#)) · checked on 7 September 2026

Data Act

Regulation (EU) 2023/2854

Applies from 12 September 2025. The Article 3(1) duty applies to connected products placed on the market after 12 September 2026.

Source: Regulation (EU) 2023/2854, Article 50 ([primary source](#)) · checked on 7 September 2026

DORA

Regulation (EU) 2022/2554

Applies from 17 January 2025.

Source: Regulation (EU) 2022/2554, Article 64 ([primary source](#)) · checked on 7 September 2026

Packaging and Packaging Waste Regulation

Regulation (EU) 2025/40

Applies from 12 August 2026. Article 67(5) applies from 12 February 2029.

Source: Regulation (EU) 2025/40, Article 79 ([primary source](#)) · checked on 7 September 2026

Critical Raw Materials Act

Regulation (EU) 2024/1252

In force. The public recycled-content statement for permanent magnets under Article 29(1) starts on the later of 24 May 2027 and two years after the calculation and verification act under Article 29(2) enters into force. That act was due by 24 May 2026 and is not published, so the start date is open.

Source: Regulation (EU) 2024/1252, Article 29(1) and (2) ([primary source](#)) · checked on 7 September 2026

Product Liability Directive, recast

Directive (EU) 2024/2853

Software is within scope. Member States are to transpose it by 9 December 2026.

Source: Directive (EU) 2024/2853, Article 22(1) ([primary source](#)) · checked on 7 September 2026

Carbon Border Adjustment Mechanism

Regulation (EU) 2023/956

Applies from 1 October 2023. Articles 5, 10, 14, 16 and 17 apply from 31 December 2024, and the definitive regime articles from 1 January 2026.

Source: Regulation (EU) 2023/956, Article 36 ([primary source](#)) · checked on 7 September 2026

Taxonomy Regulation

Regulation (EU) 2020/852

Articles 4 to 7 and Article 8(1) to (3) apply from 1 January 2022 for the first two environmental objectives and from 1 January 2023 for the other four.

Source: Regulation (EU) 2020/852, Article 27(2) ([primary source](#)) · checked on 7 September 2026

Corporate Sustainability Reporting Directive

Directive (EU) 2022/2464

Article 4 applies from 1 January 2024 for financial years starting on or after that date. Amended by Directive (EU) 2026/470, whose Articles 1, 2 and 3 are to be transposed by 19 March 2027.

Source: Directive (EU) 2022/2464, Article 7, read with Article 5 of Directive (EU) 2026/470 ([primary source](#)) · checked on 7 September 2026

Standards and frameworks

Not legislation. Referenced by legislation, or used by audits. The rows below name each item and what it is for. Where this pass did not read a standards body's own text at source, the row states no date, and that absence is deliberate.

EN 18216, EN 18219, EN 18220, EN 18221, EN 18222, EN 18223

CEN, CENELEC and ETSI

References cited in the Official Journal by Implementing Decision (EU) 2026/1736 of 14 July 2026, giving presumption of conformity for the requirements of Articles 10 and 11 of Regulation (EU) 2024/1781 that they cover.

Source: Commission Implementing Decision (EU) 2026/1736, Recitals 3 and 4, read with Article 41(2) of Regulation (EU) 2024/1781 ([primary source](#)) · checked on 7 September 2026

EN 18239 and EN 18246

CEN, CENELEC and ETSI

Access rights, and data authentication and integrity. Not among the references cited by Implementing Decision (EU) 2026/1736, so no presumption of conformity attaches to them today.

Source: Commission Implementing Decision (EU) 2026/1736, Recital 3 ([primary source](#)) · checked on 7 September 2026

Audit and assurance assessment architecture for AI

BSI, Germany

A draft criteria catalogue for assessing AI systems, machine-readable, built on an assurance engagement methodology. Not read at source in this pass, so no date is stated here.

ISO/IEC 42001

ISO and IEC

A management system standard for artificial intelligence, certifiable by a third party. COADF is an implementation-level companion to it and does not stand against it.

AI Risk Management Framework

NIST, United States

A voluntary framework with four functions, plus a profile for generative systems. Referenced here for its vocabulary, not as an obligation.

EPCIS and Digital Link

GS1

The identifier and event layer a passport is reached through. Open specifications; naming one is not a claim of any status with the body.

ECLASS and the IEC Common Data Dictionary

ECLASS and IEC

Licensed vocabularies. Treated as isolated adapters under P-7, so that a licence change never reaches the core.

Code of practice for general-purpose AI

European Commission

Voluntary. Its safety chapter addresses providers of models with systemic risk.

Monitor, do not build against

Ecodesign delegated acts by product category

Each act is adopted under Article 4 of Regulation (EU) 2024/1781 and creates its own rule set. Follow the acts, not the working plan.

Source: Regulation (EU) 2024/1781, Article 4 ([primary source](#)) · checked on 7 September 2026

A scheme for digital product passport service providers

Recital 40 of Regulation (EU) 2024/1781 says the Commission could carry out an impact assessment to investigate whether such a scheme is appropriate. Nothing further is in the text.

Source: Regulation (EU) 2024/1781, Recital 40 ([primary source](#)) · checked on 7 September 2026

Forced Labour Regulation

Regulation (EU) 2024/3015 applies from 14 December 2027, with Articles 5(3), 7, 8, 9(2), 11, 33, 35 and 37(3) applicable since 13 December 2024.

Source: Regulation (EU) 2024/3015, Article 40 ([primary source](#)) · checked on 7 September 2026

Corporate Sustainability Due Diligence Directive

Directive (EU) 2024/1760, as amended by Directive (EU) 2026/470 of 24 February 2026: transposition of the amending article by 26 July 2028, application of those measures from 26 July 2029, and the measures for Article 16 for financial years starting on or after 1 January 2030. Scope thresholds, on the principal route for EU companies: more than 5 000 employees on average and a net worldwide turnover above EUR 1 500 000 000. Other categories, including non-EU companies and the franchising and licensing route, carry their own criteria, and the figures above are not those. The amending directive removes harmonised Union conditions of civil liability, which is not the same as removing civil liability; national law continues to apply.

Source: Directive (EU) 2026/470, Article 4(2) and Article 5 ([primary source](#)) · checked on 7 September 2026



COADF · Compliance-Oriented AI Development Framework · Version 2.2 · September 2026

COADF is publicly documented as a development framework. No authority has assessed, audited or endorsed it.

Proposed licence: CC BY-SA 4.0 for the text, MIT for the schemas. Not yet granted.

COADF, by L. C. Hogrefe (AnyLAI), version 2.2, 2026

