AI REGULATION AND RISK READINESS

A Practical EU AI Act and NIST GenAI Crosswalk

30 AUGUST 2024 • ENTERPRISE AI SERIES



Paul Forrest

PAPER EIGHT · ENTERPRISE AI, 2023 TO 2024

About this paper

This is paper eight of nine on the enterprise adoption of generative artificial intelligence.

August 2024 is the point at which the subject moves from principles to implementation. A European regulation entered into force on 1 August, and voluntary risk guidance specific to generative systems was published in the United States on 26 July.

The legal chronology and the timetable are taken from the regulation as published. The crosswalk, the evidence pack and the roadmap are mine, and this paper is not legal advice, which the disclaimer overleaf says in terms.

HOW TO CITE

Forrest, P. (2024). *AI Regulation and Risk Readiness: A Practical Crosswalk Between European Law and Voluntary Risk Guidance.* Enterprise AI, 2023 to 2024, Paper Eight. Version 1.0, 30 August 2024.

HOW TO READ THE EVIDENCE TAGS

Every substantive claim carries a tag stating what kind of evidence stands behind it. A claim resting on two kinds carries both.

PRIMARY An institutional or official document, cited from the issuing body.

LEGAL A statute, regulation, executive order, legislative proposal or court decision.

AUTHOR My own framework or judgement, proposed and not yet proven.

A NOTE ON DATES

Research for this paper closed on 27 August 2024. The phased timetable in section three is taken from Article 113 as published in the Official Journal on 12 July, and if that article is amended later the dates here will not follow it.

LEGAL AND REGULATORY NOTICE

This paper describes a European regulation, United States federal guidance and other legislative instruments throughout. **Nothing in this document constitutes legal advice.** The provision of legal advice sits outside the terms of any engagement with me, and any question about whether a particular organisation is in scope, what role it occupies, or what it must do is a matter for its own qualified advisers.

Legal status is stated precisely and deliberately. Proposed, agreed, adopted, signed, published, in force and applicable are seven distinct states, and an organisation planning against the wrong one will act too early or discover a deadline late. Where I draw a practical implication that an instrument does not itself state, the text says so.

The crosswalk in section five maps a legal obligation against a voluntary practice. It does not assert that the two are equivalent, and satisfying the second does not discharge the first.

Contents

FRONT MATTER

	Abstract and summary of the argument	4

THE PAPER

1	Two instruments of entirely different kinds	5
2	How the regulation was made, step by step	7
3	What applies on 30 August 2024, and what does not	9
4	The voluntary profile, and what it is good for	11
5	Inventory, roles and classification	12
6	The evidence pack	14
7	A roadmap to February and beyond	16

END MATTER

	Method, evidence and limits	18
	Sources considered and set aside	19
	About the author	20
	References	21
	Further sources consulted	22

ABSTRACT AND SUMMARY OF THE ARGUMENT

In force since 1 August, applicable to nobody yet, and the distinction is the whole paper

The European artificial intelligence regulation entered into force on 1 August 2024, 20 days after its publication in the Official Journal.[1] LEGAL On 30 August 2024 not one of its substantive provisions has become applicable, and the earliest application date is 2 February 2025.[1] LEGAL Entry into force and applicability are different things. Treating them as the same is the costliest error in this subject.

Six days before that, the United States standards institute published a profile addressing generative artificial intelligence risks as a companion to its existing risk management framework.[2] PRIMARY It is voluntary, cross-sectoral, and organised around risks and suggested actions. It imposes nothing.

This paper builds a working crosswalk between the two without claiming they are equivalent. One creates legal duties on a fixed timetable; the other offers a method. What makes the crosswalk worth doing is that a single evidence system can serve both, and most of the work an organisation must do in the next 18 months is evidence work it would benefit from regardless. AUTHOR

WHAT THIS PAPER ARGUES

One evidence system can support multiple governance regimes. Legal applicability and voluntary practice must stay separate in the record, because a regulator asking whether an obligation was met will not accept a voluntary framework as the answer.

- The chronology matters. The Parliament adopted its position on 13 March 2024, the Council adopted the regulation on 21 May, the presidents signed it on 13 June, it was published on 12 July and it entered into force on 1 August.[1] LEGAL The date of 13 June is signature, and describing it as adoption is a common and material error.
- Nothing applies yet. Prohibitions and the artificial intelligence literacy duty apply from 2 February 2025, general-purpose model obligations and governance from 2 August 2025, general application from 2 August 2026, and one further category from 2 August 2027.[1] LEGAL
- Role determines duty. Provider, deployer, importer and distributor carry different obligations, and an organisation that fine-tunes or rebrands a model may find it occupies a different role from the one it assumed. LEGAL
- The voluntary profile is a method rather than a shield. It identifies a set of generative risks with suggested actions against the four functions of the parent framework, and following it produces useful evidence and discharges no legal obligation.[2] PRIMARY
- The inventory should be built now. Every route through the next two years starts with knowing which systems exist, what they do, where they operate and who owns them, and an organisation without that record has 23 months to build one and will need most of them.

SECTION 1

Two instruments of entirely different kinds

One creates legal duties on a fixed timetable. The other offers a method. A crosswalk between them is useful and an equivalence between them is false.

Compliance programmes in this area are being built against a mixture of sources, and the mixture obscures a distinction that determines what happens when somebody asks a question. A regulation confers obligations, timetables and penalties. A voluntary profile suggests actions. The first is enforceable and the second is not, and no amount of adherence to the second answers a question posed under the first.

The crosswalk in this paper exists because the underlying work overlaps heavily. Both instruments want an organisation to know what its systems do, to have tested them, to have designed human oversight, to log, to monitor and to keep records. An organisation building that evidence once can present it in two directions, and the saving is substantial. What must not happen is that the two records merge, because a regulator asking whether a legal duty was discharged will not accept a voluntary framework as the answer. AUTHOR

A CROSSWALK, NOT AN EQUIVALENCE



FIGURE 1 Legal obligations on the left, voluntary practice on the right, and a shared evidence layer between them. The arrows run only one way, because voluntary practice can produce evidence for a legal duty and cannot satisfy it.

My framework. AUTHOR

What a crosswalk can and cannot do

1. It can identify where one piece of evidence serves two purposes, which is most of the documentation work.
2. It can sequence effort, so that work required by the earliest legal deadline is done first and voluntary practice fills the gaps.
3. It cannot map obligations onto suggestions. Where the regulation requires something the profile does not address, the requirement stands alone.
4. It cannot be used as a compliance argument. The correct use is internal planning, and the correct external position is the legal analysis an organisation's advisers produce.

SECTION 2

How the regulation was made, step by step

The regulation took eight steps over three years, and naming each one correctly matters because the wrong term produces the wrong planning date.

The chronology below is set out in full because summaries of it circulate with errors, and because two of the steps are routinely conflated in ways that shift planning assumptions by months. The steps below are taken from the records of the institutions concerned, and the last three of them from the regulation itself.[1,7,8] LEGAL / PRIMARY

THE LEGISLATIVE CHAIN, WITH THE CORRECT TERM FOR EACH STEP

DATE	WHAT HAPPENED	WHY THE DISTINCTION MATTERS
21 April 2021	Commission proposal	A legislative proposal creates no obligation. Papers written between 2021 and 2023 that describe the Act as forthcoming law were describing a negotiating position.
6 December 2022	Council general approach	The Council agrees its own negotiating position. The text at this point differs materially from what was finally adopted.
9 December 2023	Provisional political agreement	A trilogue agreement between negotiators. It has no legal effect and the text is not final, which is why the December 2023 coverage describing settled obligations was premature.
13 March 2024	European Parliament adopts first-reading position	By 523 votes to 46 with 49 abstentions. A genuine adoption by one institution, and not the completion of the procedure.
21 May 2024	Council formally adopts	This completes the legislative procedure. If a single date is wanted for when the law was made, this is the better candidate.
13 June 2024	Signature by the two presidents	The date the act bears, which is why the citation reads of 13 June 2024. It is a signature and not an adoption, and this is the most commonly mis-stated date in the whole chain.
12 July 2024	Publication in the Official Journal	The regulation acquires its citation number on publication. Any document citing that number before this date is citing something that did not exist.
1 August 2024	Entry into force	Entry into force came 20 days after publication. The instrument now exists in law, and no substantive obligation has yet become applicable.

The error to avoid

The date of 13 June 2024 is frequently described as the date the regulation was adopted. It is the date the presidents of the Parliament and the Council signed it, which is why the act bears that date in its own citation.[1] LEGAL Adoption by the Parliament was 13 March and adoption by the Council was 21 May. An organisation that takes 13 June as the adoption date will not derive a wrong deadline, because the staged application dates are fixed calendar dates written into the text. It will still look imprecise in front of anyone who knows the file.

A second point of care concerns the citation number. The number is assigned on publication in the Official Journal, which happened on 12 July 2024. Any document dated before that date which cites the number is citing something that did not then exist. LEGAL

SECTION 3

What applies on 30 August 2024, and what does not

Nothing applies. The instrument is in force and its earliest substantive obligation begins on 2 February 2025.

The regulation entered into force on 1 August 2024, being the twentieth day following its publication.[1] LEGAL Entry into force means the instrument exists as law. It does not mean any duty it contains has begun to bite, and the regulation itself sets out a staged calendar in Article 113.

WHAT APPLIES, AND WHEN



FIGURE 2 The phased timetable from the regulation as published. On the date of this paper every band sits in the future, which is the fact about the instrument most often misunderstood.

Article 113 of the regulation as published on 12 July 2024. LEGAL

THE PHASED TIMETABLE, FROM ARTICLE 113 AS PUBLISHED

FROM	WHAT BECOMES APPLICABLE	WHAT AN ORGANISATION SHOULD BE DOING NOW
2 February 2025	Chapter I, general provisions including the artificial intelligence literacy duty, and Chapter II, prohibited practices.	Screen the inventory against the prohibited categories, and establish what literacy means for staff who deal with these systems. Five months remain at the date of this paper.
2 August 2025	Chapter III Section 4, Chapter V on general-purpose models, Chapter VII on governance and Chapter XII on penalties, together with Article 78, with the exception of Article 101.	Establish whether the organisation is a provider of a general-purpose model, which fine-tuning and rebranding can cause, and open supplier conversations about documentation.
2 August 2026	General application, including the high-risk obligations attaching to the Annex III categories and the transparency duties in Article 50.	Classification, conformity work, technical documentation, human oversight design and post-market monitoring. This is the substantial programme and it takes most of two years.
2 August 2027	Article 6(1) and its corresponding obligations, covering systems that are safety components of products already regulated under Union sectoral legislation.	Relevant to manufacturers of regulated products. Align with the existing conformity assessment route rather than building a parallel one.

Article 111 adds transitional provisions that are frequently omitted from summaries and matter to organisations with existing estates. General-purpose models placed on the market before 2 August 2025 have until 2 August 2027 to comply. High-risk systems placed on the market before 2 August 2026 are caught only where they undergo significant changes in design after that date, with a later date applying to public authority deployers. Systems forming components of certain large-scale information technology systems have until the end of 2030.[1] LEGAL

THE SENTENCE TO STRIKE FROM ANY BOARD PAPER WRITTEN THIS MONTH

Anything asserting that the regulation now applies, that prohibitions are in effect, or that the organisation must already comply. None of those is true on 30 August 2024, and a board told otherwise will fund the wrong work on the wrong timetable. LEGAL

The accurate formulation is that the regulation is in force, that its first obligations apply from 2 February 2025, and that the substantial high-risk programme applies from 2 August 2026. Stating it that way also makes the case for starting now, because 2026 is closer than it sounds once classification and conformity work are scoped.

Why the first deadline is easier than it looks and the second is not

February 2025 brings prohibited practices and an artificial intelligence literacy duty. Most organisations will find that they operate nothing in the prohibited categories, so the work is a screening exercise against the inventory. The literacy duty is a training and record-keeping matter which the fifth and seventh papers in this series already point at.

The August 2026 deadline is a different order of undertaking wherever high-risk systems exist. Classification, technical documentation, risk management across the lifecycle, data governance, logging, human oversight design, accuracy and robustness, conformity assessment and post-market monitoring together constitute a programme rather than a project. At the date of this paper 23 months remain, and the first six of them will be spent establishing which systems are in scope. AUTHOR

SECTION 4

The voluntary profile, and what it is good for

Published 26 July, cross-sectoral, and organised around generative risks with suggested actions. It creates no obligation and supplies something the regulation does not.

The United States standards institute published a generative artificial intelligence profile on 26 July 2024, as a companion resource to the risk management framework it published in January 2023.[2] PRIMARY It identifies risks that are unique to or exacerbated by generative systems and sets out a substantial number of suggested actions organised against the four functions of the parent framework.

It was produced under the executive order of October 2023, it is voluntary, and it is cross-sectoral.[2] PRIMARY Three further publications and one software release appeared the same day. Two of those publications are final, and the third, addressing misuse risk for dual-use foundation models, was issued as an initial public draft.[3,4] PRIMARY That distinction matters in a compliance record and is easily lost.

What it supplies that the regulation does not

A regulation states what must be achieved. It does not, and should not, tell an organisation how to test a generative system for a given risk. The profile is a catalogue of things an organisation might do, from which a proportionate set can be selected. Used that way it is the more practically useful of the two documents for anybody who has to build an evaluation capability this year.

SPECIFICATION

How to use a voluntary profile without misusing it

- Select from it proportionately. It is a catalogue rather than a checklist, and an organisation attempting every suggested action has misread the instrument.
- Record why each selection was made and each omission accepted, because the reasoning is the evidence.
- Keep the record separate from any legal compliance record, so that the two cannot be confused by a later reader.
- Never describe adherence to it as compliance with anything. It has no compliance state and no certification attaches to it.

One further instrument belongs in the same category and is often overlooked. A management system standard for artificial intelligence was published in December 2023 and is certifiable, which distinguishes it from both documents above.[5] PRIMARY Certification against it demonstrates that a management system exists and operates. It does not demonstrate that any individual system is safe, and it discharges no legal obligation.

SECTION 5

Inventory, roles and classification

Every route through the next two years starts with a list. Most organisations do not have one and will need six months to build it.

The first question under either instrument is what systems the organisation has. The second is what role it occupies in relation to each. Both are administrative, and both take far longer than anybody expects because the information sits in a dozen places and some of it sits nowhere.

SPECIFICATION

The minimum inventory record, per system

- A name, an owner and the business process it supports.
- Its purpose, stated in terms of what it does.
- The geography in which it operates and the people it affects.
- The supplier, the model and the version, with the contract reference.
- The role the organisation occupies in relation to it.
- Whether a human decides, reviews or is absent from the decision.
- The data it touches and the classification of that data.
- The date of the last evaluation and by whom.

FOUR ROLES, AND HOW ORGANISATIONS ACQUIRE THEM WITHOUT NOTICING

ROLE	BROADLY WHO THIS IS	HOW IT IS ACQUIRED UNINTENTIONALLY
Provider	Develops a system or a general-purpose model, or has one developed, and places it on the market or puts it into service under its own name or trademark.	By putting your own name on a purchased system, or by modifying a system substantially. This is the role organisations most often occupy without having considered it.
Deployer	Uses a system under its own authority, other than in a personal non-professional activity.	The default position for most enterprises, and the one most internal assessments assume without testing the assumption.
Importer	Places on the Union market a system bearing the name or trademark of a person established outside the Union.	By procuring from a non-Union supplier and making the system available within the Union.
Distributor	Makes a system available on the Union market without being the provider or importer.	By reselling or embedding a third-party system in an offering to customers.

The role most often misassigned

Organisations assume they are deployers, because they bought something and are using it. Two common activities disturb that assumption. Putting the organisation's own name or trademark on a system, and modifying a system substantially, can both move an entity into the provider role, which carries a materially heavier set of obligations. LEGAL Whether any particular arrangement crosses that line is a legal question for advisers, and this paper does not answer it.

The practical instruction is narrower and safe to give. Where an organisation fine-tunes a model, rebrands a purchased system, or embeds a third-party system into something it sells, those facts should be recorded in the inventory now, so that the legal analysis when it happens has something to work from. An inventory that records only the supplier name will not support that analysis.

Classification, and the sequence that saves effort

Screen first against the prohibited categories, whose deadline is nearest. Then identify systems that may fall within the high-risk categories, which is the work that determines the size of the 2026 programme. Then consider the transparency duties, which attach to particular kinds of interaction and content whatever the risk level. Running these three in the other order is the most common way to spend a quarter and learn nothing. AUTHOR

SECTION 6

The evidence pack

Each system needs nine records. Most organisations hold fragments of all nine and an assembled set of none.

Both instruments, and every sectoral regulator likely to ask, want broadly the same things. The pack below is organised so that one set of records serves whichever enquiry arrives, and so that the work is done once.

THE EVIDENCE PACK

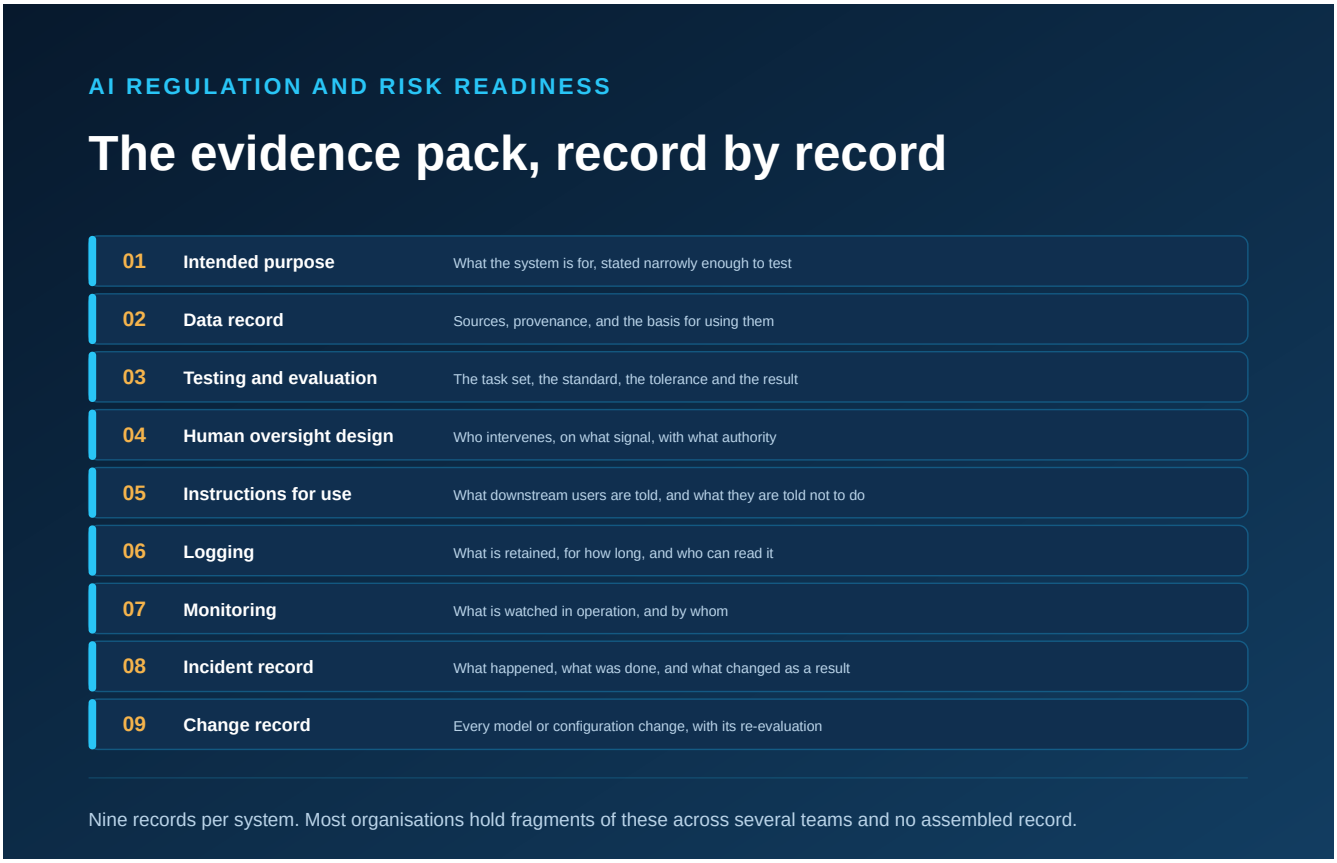


FIGURE 3 Nine records per system. Most organisations hold fragments of these across several teams and no single assembled record, which is what makes an enquiry expensive.

My framework, mapped against the documentation themes of both instruments. AUTHOR / LEGAL

THE EVIDENCE PACK, RECORD BY RECORD

RECORD	WHAT IT CONTAINS	WHICH REGIME IT SERVES
Intended purpose	What the system is for, what it is not for, the population affected, and the decisions it influences.	Both. It is the first question under the regulation and the framing question in the voluntary profile.
Data record	Sources, lawful basis, quality measures, known gaps and representativeness for the intended population.	Both, and it overlaps substantially with existing data protection documentation.
Testing and evaluation	Method, sample, results with limitations, and the standard against which adequacy was judged.	Both. This is the record most organisations cannot produce today.
Human oversight design	What the reviewer sees, what authority they have, what training they received, and evidence the oversight operates.	Both. Oversight asserted but not evidenced is the most common weakness in current documentation.
Instructions for use	What a deployer is told, including limitations, conditions and foreseeable misuse.	Primarily the legal regime, and primarily relevant where the organisation is a provider.
Logging	What is recorded, for how long, who may access it, and whether it is sufficient to reconstruct a disputed output.	Both, and it is also the practical precondition for investigating an incident.
Monitoring	What is watched in production, by whom, on what cadence, and what triggers a review.	Both.
Incident record	What happened, severity, response, remediation and any external notification.	Both, with near misses recorded separately because they arrive earlier.
Change record	Model, prompt, corpus and configuration changes with dates, and the re-evaluation that followed each.	Both. Without it, every other record describes a system that may no longer exist.

The two records that are almost always missing

Testing and evaluation is the first. Organisations have testing for conventional software and have not built an equivalent for probabilistic systems, and the absence is usually discovered when somebody asks how they know the thing works. The honest answer in most cases this year is that a small number of people tried it and were satisfied, which is a legitimate starting point and not a record.

The change record is the second, and its absence quietly invalidates everything else. Where a supplier updates a model without notice, an evaluation performed in March describes a system that no longer exists, and an organisation that cannot say when the behaviour changed cannot say what its evidence covers. That is the practical argument for negotiating change notice into supplier contracts, which the fifth paper in this series set out. AUTHOR

> ON DOCUMENTATION THAT DESCRIBES AN INTENTION
>
> Human oversight is the record most often present and least often evidenced. A statement that a qualified person reviews the output is a description of a design. What is wanted is evidence that the review happens, that the reviewer can see enough to review, that they have authority to reject, and that rejection has occurred. The last of those turns an assertion into a record.

SECTION 7

A roadmap to February and beyond

The work falls into four blocks, and the legal calendar sets their order.

What follows assumes an organisation with European exposure, a portfolio of assisted workflows and no current inventory, which describes a large share of the market at the end of August 2024.

Now to November 2024, establish the facts

Build the inventory. Determine the organisation's role for each system, recording the facts. Conclusions come later and from advisers. Screen against the prohibited categories, which is a short exercise once the inventory exists. Open supplier conversations about documentation, change notice and testing access, because supplier timelines are longer than internal ones and the requests take a quarter to land.

December 2024 to February 2025, the first deadline

Complete the prohibited-practice screen and record each conclusion. Establish what artificial intelligence literacy means for the organisation's own staff and record how it is delivered. Both of these are modest pieces of work, and both need to be evidenced, because doing them without a record proves nothing.

March 2025 to August 2025, the model layer

Determine whether the organisation provides a general-purpose model, which fine-tuning can cause and which most organisations have not considered. Obtain supplier documentation for the models in use. Begin the governance arrangements that become applicable in August.

September 2025 to August 2026, the substantial programme

Classification of systems that may be high risk, technical documentation, lifecycle risk management, data governance, logging, human oversight design, accuracy and robustness work, conformity assessment and post-market monitoring. This is where the cost sits, and an organisation that reaches September 2025 without knowing which systems are in scope has roughly 11 months to do two years of work.

> QUARTERLY REASSESSMENT, AND WHY IT IS NOT OPTIONAL
>
> Three things will change during this programme. Harmonised standards and Commission guidance will appear and will alter what compliance looks like in detail. Suppliers will change models, which invalidates evaluations. And the organisation's own portfolio will grow.
>
> A compliance position set once and filed is a position that decays. Reassess quarterly against the inventory, with legal review at each point where the organisation's role or a system's classification might have moved.
>
> One of those processes has already begun. On 30 July 2024 the Commission opened a consultation and a call for participants on the first code of practice for general-purpose models, with the drafting itself to follow. There is no text to plan against yet, and an organisation that supplies or substantially modifies such a model should be watching it. PRIMARY

One caution about other jurisdictions. Other instruments are moving on their own timetables, and one American state has already enacted a statute of its own.[6] LEGAL An organisation building a European programme should build it so that the evidence is portable, because the probability that Europe remains the

only regime requiring this documentation is low.

END MATTER

Method, evidence and limits

How the sources in this paper were chosen and dated, and what the reader should distrust.

Sources, and what was excluded on principle

Primary legal and institutional instruments only, cited from their issuing bodies. No commentary, law firm briefing or consultancy interpretation is cited, because in this subject the secondary literature is where the errors of legal status originate.

The Article 113 problem

The phased timetable in section three is taken from the text of the regulation as published in the Official Journal on 12 July 2024. This is stated because the text of a regulation can be amended by later instruments, and a reader consulting a live web version at a future date may find dates that differ from those that applied in August 2024. Where that happens, the published version is what this paper describes.

Legal status

Seven states are distinguished throughout, being proposed, agreed, adopted, signed, published, in force and applicable. The paper states that nothing was applicable on 30 August 2024, which is the sentence in it that matters most and the one most often wrong elsewhere.

What is not yet knowable

Harmonised standards under the regulation had not been published at this date, so what technical compliance will require in detail is not yet knowable. No Commission guidance on the general application provisions existed. And no enforcement practice exists, because nothing is yet enforceable.

The parts that are mine

The crosswalk structure, the nine-record evidence pack and the roadmap sequencing are my own. Statements about how an organisation acquires a particular role are summaries of definitions and are not legal advice, and any organisation whose position turns on that question should obtain advice.

My interest, and my limits

I advise on artificial intelligence governance and could be a supplier of the inventory and evidence work described here. I am not a lawyer, I do not give legal advice, and the paper says so at the front.

END MATTER

Sources considered and set aside

Material I read and did not cite, mostly because it carries no legal status.

SOURCE	PUBLISHED	WHY IT WAS SET ASIDE
European Commission, questions and answers on the artificial intelligence act	Updated August 2024	Informal guidance. Useful for orientation and without legal status, so nothing here rests on it.
Organisation for Economic Cooperation and Development, Recommendation on Artificial Intelligence, as amended	3 May 2024	Principles adopted by governments. They bind nobody at the level of an organisation, and this paper is about obligations.
Law firm client alerts on the regulation	July and August 2024	I read several and cited none. This is the layer of the literature where the confusion between adoption, signature and entry into force originates, and it would be circular to rely on it.
Department for Science, Innovation and Technology, response to the consultation on the regulation white paper	6 February 2024	The United Kingdom position, which is a sector-led approach with no horizontal statute. It sits outside a paper about a European regulation, and readers with United Kingdom exposure should not read this paper as covering it.
European Commission, consultation on the first general-purpose code of practice	30 July 2024	Opened for responses on 30 July, with the drafting to follow. There is no text yet, and section seven says so.

END MATTER

About the author

Paul Forrest advises boards and executive teams on the adoption of new technology in operating businesses, with a practice grounded in delivery rather than in advisory work alone.

His background is in business process reengineering, which he has practised since the 1990s, and in applied natural language processing, which he has worked on since 2014. He writes about regulation as a board adviser rather than as a lawyer, which means describing what an instrument says, how firmly it says it, and what follows for an organisation, while leaving the question of what the law requires of a particular company to that company's advisers.

He writes on the condition that a claim carries its evidence, and this series is an attempt to hold that line while a subject moves quickly.

END MATTER

References

Eight works cited in the text, each with its original date of publication. None is later than 27 August 2024.

1. **European Union.** Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence. Published in the Official Journal on 12 July 2024 and entered into force on 1 August 2024. Article 113 sets the phased application timetable and Article 111 the transitional provisions.
   PUBLISHED 12 JULY 2024

2. **National Institute of Standards and Technology.** *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile.* NIST AI 600-1. Voluntary, cross-sectoral, produced under Executive Order 14110.
   PUBLISHED 26 JULY 2024

3. **National Institute of Standards and Technology.** *Secure Software Development Practices for Generative AI and Dual-Use Foundation Models, An SSDF Community Profile.* NIST SP 800-218A. Final publication.
   PUBLISHED 26 JULY 2024

4. **National Institute of Standards and Technology.** *Managing Misuse Risk for Dual-Use Foundation Models.* NIST AI 800-1. Initial public draft, cited as a draft.
   PUBLISHED 26 JULY 2024

5. **International Organization for Standardization and International Electrotechnical Commission.** *ISO/IEC 42001, Information technology, Artificial intelligence, Management system.* Certifiable, unlike either instrument above.
   PUBLISHED 18 DECEMBER 2023

6. **Colorado General Assembly.** Senate Bill 24-205, concerning consumer protections in interactions with artificial intelligence systems, signed 17 May 2024.
   PUBLISHED 17 MAY 2024

7. **European Parliament.** Position adopted at first reading, Texts Adopted P9_TA(2024)0138.
   PUBLISHED 13 MARCH 2024

8. **Council of the European Union.** Formal adoption of the artificial intelligence regulation.
   PUBLISHED 21 MAY 2024

END MATTER

Further sources consulted

Six sources I read for this paper and did not cite. They are listed so that the reading behind it can be seen in full, and none is later than 27 August 2024.

- **European Commission.** Decision establishing the European AI Office, applying from 21 February 2024.
  PUBLISHED 24 JANUARY 2024
- **Council of Europe.** Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, adopted by the Committee of Ministers. Cited as adopted, since it has not been opened for signature at the date of this paper.
  PUBLISHED 17 MAY 2024
- **National Institute of Standards and Technology.** *Artificial Intelligence Risk Management Framework 1.0.* NIST AI 100-1, the parent framework.
  PUBLISHED 26 JANUARY 2023
- **The White House.** *Executive Order 14110.* The instrument under which the generative profile was produced.
  PUBLISHED 30 OCTOBER 2023
- **AI Seoul Summit.** *Seoul Declaration* and the frontier artificial intelligence safety commitments.
  PUBLISHED 21 MAY 2024
- **National Institute of Standards and Technology.** *A Plan for Global Engagement on AI Standards.* NIST AI 100-5.
  PUBLISHED 26 JULY 2024

WHERE THIS GOES NEXT

The final paper takes systems that act rather than answer

This paper concerned documentation for systems that produce content. The ninth, and the last in the series, will concern systems that select tools and take multi-step actions, which changes the control question from reviewing output to governing permissions. It is due in December.

The vendors are already describing such systems, and the figures attached to them are the vendors' own. I expect that paper to be more cautious than the market coverage, because the published capability evidence is thin now and will probably still be thin in December.

THREE RECORDS THE FINAL PAPER DEPENDS ON

- An inventory recording role, geography, supplier, model version and human involvement for every system.
- A written prohibited-practice screen with a conclusion for each entry.
- A change record, so that every other piece of evidence can be tied to a version of the system it actually describes.

Paul Forrest · 30 August 2024 · Version 1.0 · Research closed 27 August 2024