


GOVERNANCE / OCTOBER 2026

What a regulator actually asks

A management reading of selected EU, US credit and UK requirements and guidance



THE DECISION IN BRIEF

A defensible consequential decision connects its authority, inputs, tests and remedy in a record the organization can retrieve and challenge.

- 01

Begin with the actual jurisdiction, activity and organizational role. Selected EU, US credit and UK examples show why a generic AI compliance checklist cannot settle applicability.
- 02

Connect purpose, authority, provenance, validation, intervention and remedy. This mayden decision-evidence chain is a management framework; legal duties must be mapped separately to the relevant use.
- 03

Test the chain on consequential decisions, including difficult cases and appeals. Fund the missing links that prevent reconstruction, intervention or correction, and verify that the record survives supplier change.

THE MAYDEN VIEW


A defensible decision has a record that can be retrieved, challenged and corrected.

INSIDE THIS PERSPECTIVE

The argument at a glance	3
01 / Identify the decision and applicable duty	4-5
Exhibit 01 / Match evidence to the jurisdiction, role and instrument	6
02 / Connect claims to inspectable evidence	7
Exhibit 02 / Different claims require different proof	8
03 / Test intervention and expose missing links	9-10
04 / Assign an evidence programme with a decision	11
The leadership agenda	12
Method and references	13

ONE DECISION, FOUR CONNECTED QUESTIONS

Make the claim inspectable



EVIDENCE TO REVIEW RECORD / CONCEPTUAL ARTWORK

01 Identify the decision and applicable duty

Reconstruct the decision a customer can challenge

02 Connect claims to inspectable evidence

Connect the evidence to the consequential action

03 Test intervention and expose missing links

Test the claim, the reviewer and the remedy

04 Assign an evidence programme with a decision

Make the evidence portable when a supplier changes

A defensible decision has a record that can be retrieved, challenged and corrected.

CHAPTER 01

01

Identify the decision and applicable duty

Reconstruct the decision a customer can challenge

A customer challenges an AI-assisted refusal. Can the organization retrieve the information used, identify the authority behind the decision and demonstrate a working route to correction? This illustrative test crosses legal, technical and operating responsibilities. Separate policy folders cannot answer it unless the records connect to the actual decision.

Design traceability into the service. Link purpose, deployed version, input record, permitted action and subsequent review. The organization need not preserve every internal model computation; it needs reliable evidence at the level where responsibility sits. An

explanation reconstructed without the actual inputs is a hypothesis about the decision, not its record.

The sources below inform this framework without establishing it as a universal requirement. The NIST AI Risk Management Framework is explicitly voluntary and intended for application across the AI lifecycle.^[1] Its value here is a management vocabulary. The legal duties applicable to an actual organization still depend on its activities, role, market and use of information. The board should insist that these distinctions remain visible in its reporting.

CHAPTER 01

01

Identify the decision and applicable duty

A jurisdiction map is more useful than a universal checklist

The EU AI Act distinguishes providers, deployers and general-purpose model providers, with obligations shaped by the system and its use. ^[2]

A company buying a model is therefore not automatically subject to every duty of a model developer. Equally, purchasing a service does not settle responsibility for the business decisions made with it. The first evidence artefact should name the organization's role, the relevant system boundary, the countries involved and the specific obligation being addressed.

Dates also matter. The Commission's current service desk states that transparency obligations apply from 2 August 2026, Annex III high-risk rules from 2 December 2027 and relevant product-embedded high-risk rules from 2 August 2028. ^[3] Providers of systems marketed before 2 August 2026 have until 2 December 2026 for the Article 50(2) marking and detection obligation. Other transitions also

require checking. An undated compliance slide that puts all AI obligations on one milestone is unsuitable evidence. Legal owners should confirm the operative text and applicable exceptions before a deployment decision; this report is a management analysis as of 10 October 2026.

Outside that regime, existing sector duties can be more directly relevant than an AI-specific label. In US credit, Regulation B requires the reasons for covered adverse actions to identify the specific principal grounds, with notice and request procedures depending on the circumstances. ^[4] In the UK, the ICO provides separate guidance on data protection and explaining AI-assisted decisions. ^[5] Neither example establishes a global requirement. They illustrate why the evidence inventory must begin with an actual activity rather than a generic list of frameworks. The official interpretation of Regulation B also requires disclosed reasons to relate to factors actually considered or scored. ^[4] A fluent explanation that substitutes a more plausible reason fails that evidential test.

SCOPED LEGAL AND GUIDANCE COMPARISON

Match evidence to the jurisdiction, role and instrument

Selected instruments, checked 10 October 2026

EU system: provider or deployer	Instrument and status AI Act; duties and dates depend on role and use	Evidence to bring Role/classification assessment, operative text and system records
US creditor: covered adverse action	Instrument and status Regulation B, section 1002.9; binding within its scope	Evidence to bring Notice/request route and principal reasons tied to actual factors
UK personal-data processing	Instrument and status Data-protection law, interpreted through ICO guidance	Evidence to bring Purpose, data handling and explanation assessment
Management assurance design	Instrument and status NIST AI RMF and UK assurance guidance; not universal legal duties	Evidence to bring Tests of claimed control, intervention and remedy
THE MANAGEMENT IMPLICATION A control record should answer a specific obligation or management claim; documents cannot repair an applicability error.		

Commission guidance gives 2 August 2026 for transparency, 2 December 2027 for Annex III high-risk rules and 2 August 2028 for relevant product-embedded systems. Other transitions and exceptions apply. Confirm operative law for the actual use. Sources: [1], [2], [3], [4], [5], [6]

CHAPTER 02

Connect claims to inspectable evidence 02

Connect the evidence to the consequential action

mayden proposes six connected records. Purpose identifies the decision and intended benefit. Authority identifies who may approve, execute, override and stop it. Provenance records sources, rights and transformations. Validation holds relevant tests and limitations. Intervention shows how the action is constrained. Remedy follows an error to correction and communication. These are management design choices, not identical forms required by every regulator.

A purchasing assistant that drafts a contract summary needs source clauses and the buyer's review. Authority to sign or alter terms additionally requires a delegation boundary, transaction limits, reliable identity and a recovery process. A model card can inform component assessment; it cannot establish permission for a particular purchase. Keep authority and quality in the same decision record.

DSIT's assurance guidance distinguishes techniques used to evaluate systems and communicate justified confidence. ^[6] The practical implication is to choose evidence for the claim being made. A penetration test addresses a different question from a representative accuracy evaluation. A policy attestation addresses a different question from a demonstrated recovery drill. Keeping those distinctions prevents a stack of individually valid documents from being presented as proof of a proposition none of them tested.

Evidence quality depends on the operating boundary

Retain the deployed configuration, decision time, source snapshot or secure reference, and permissions in force. Apply deliberate retention and access rules to personal information. More logs can add exposure without making a decision easier to explain. Capture sufficient evidence for the defined claim under applicable legal requirements.

External suppliers introduce a particularly important boundary. Contractual assurances should be translated into operational capabilities: notification of material changes, access to relevant test information, incident cooperation, restrictions on data reuse and an exit path. The supplier may not disclose model internals. The buyer can still define the permitted use, test the end-to-end service, monitor outcomes and stop deployment if the evidence no longer supports the decision. Unavailable information is a limitation to manage, not an invitation to fill the gap with marketing language.

The NCSC's secure-development guidance covers design, development, deployment and operation. ^[7] Our extension is to preserve continuity between those phases in the decision record. An initial approval expires in practical terms when a materially different tool, data source or permission is introduced. A change record should state what changed and which earlier evidence remains relevant. This avoids both extremes: treating every minor update as a complete reapproval or treating a major behavioral change as routine maintenance.

ORIGINAL MAYDEN FRAMEWORK · ILLUSTRATIVE CONTROL DESIGN

Different claims require different proof

Illustrative control design

Reasons are accurate	Appropriate demonstration Compare reasons with actual decision inputs	Insufficient substitute Fluent explanation
Authority is bounded	Appropriate demonstration Attempt out-of-scope actions	Insufficient substitute Written acceptable-use policy
People can intervene	Appropriate demonstration Exercise an override under realistic load	Insufficient substitute A human named in a process chart
Errors are recoverable	Appropriate demonstration Trace and correct an affected transaction	Insufficient substitute A ticketing workflow exists
THE MANAGEMENT IMPLICATION Evidence should address the proposition actually being asserted.		

mayden analysis; examples are not observed company findings. Sources: [4], [6]

CHAPTER 03

03

Test intervention and expose missing links

Test the claim, the reviewer and the remedy

Validation should start with the claim management wants to make. If the claim is that a customer explanation is accurate, reviewers should compare it with the actual reason for the decision. If the claim is that an agent cannot exceed a spending limit, a test should attempt to cross the limit through permitted workflows and plausible misuse. If the claim is that a human catches material errors, an evaluation must include the human's workload, authority and access to evidence. A model benchmark alone cannot establish any of these operating claims.

For high-risk AI systems within its scope, Article 26 of the EU AI Act assigns oversight to people with the required competence, training, authority and support. It also addresses monitoring and retention of automatically generated logs under the deployer's control, with a six-month minimum subject to the stated legal qualifications. ^[8] These provisions require an applicability assessment; they are not a universal checklist for every company using AI.

Trace a live case under appropriate access permissions. Identify who could intervene, show the evidence available to that person and demonstrate the effect on the downstream

action. Retrieve the record through the actual retention process. This tests a specific operating claim; passing it does not certify overall legal compliance.

Sampling deserves explicit treatment. A demonstration on easy cases says little about missing documents, ambiguous requests, conflicting records or underrepresented languages. Start with the distribution of real decisions, identify where the harm could concentrate and deliberately include those cases. Keep ordinary operating performance separate from adversarial tests; both are useful, but their denominators mean different things. Do not convert a handful of successful demonstrations into an enterprise-wide reliability percentage.

Remedy is often the neglected test. An illustrative drill should introduce an incorrect input, follow its influence into a decision, suspend further affected actions and verify correction across connected systems. Verify that a responsible person can contain the consequence and explain what remains unresolved after the ticket is created. Such drills provide management with a practical view of recoverability without pretending to certify that future errors are impossible.

CHAPTER 03

03

Test intervention and expose missing links

A board review should expose gaps rather than average them away

mayden recommends a small decision sample at each material governance review. Select a routine case, a difficult case and a case where a person overrode or appealed the system. Follow each from input to outcome. Record missing links by owner and business consequence. Avoid a single maturity score that lets strong documentation compensate numerically for an absent stop mechanism. Some weaknesses are blockers; others justify a limited release while evidence improves.

The dashboard should distinguish available evidence from credible evidence. A validation document may exist but test an obsolete version. A signed policy may exist but exclude a newly connected tool. An incident register may

show few events because staff do not know how to report them. These observations do not prove negligence; they identify where management's confidence exceeds its information. A useful board discussion asks what could disprove the release case and who has the independence and standing to raise that challenge.

There is a legitimate counterargument: evidence work can become expensive bureaucracy. The answer is proportionality and reuse. A low-consequence drafting tool should not inherit the complete record set of a consequential automated decision. Common platform controls can be evidenced once and referenced by multiple uses, while decisions with different harms retain distinct tests. The test of proportionality is whether removing an artefact would weaken a material decision claim, not whether another company happens to collect it.

CHAPTER 04

04

Assign an evidence programme with a decision

Make the evidence portable when a supplier changes

The record should survive a supplier change. If a complaints service changes model version during an investigation, running the case again does not establish what the customer originally saw. Preserve the delivered answer, relevant source record, effective policy and intervention history under the applicable retention rules. Reconstructing a recorded decision is different from forcing a probabilistic model to repeat its words.

Before signing a supplier contract, ask the investigator and the service owner to specify the export they would need. Include stable record identifiers, timestamps, relevant configuration and the identities or roles responsible for consequential actions. Require an explanation of what will remain unavailable. Test the export in a format that a replacement team can read. An elaborate dashboard is weak protection if access ends with the subscription or if its records cannot be reconciled with the enterprise's own transactions.

There is a cost trade-off. Preserving every interaction indefinitely would increase storage, privacy and discovery burdens. Our proposed design begins with the decision claim and retains the smallest sufficient record, with access separation and deletion rules. Legal specialists determine applicable duties; engineering teams establish what can be captured accurately; operations verify that the record answers real questions. Assign funding for this work in the service budget. Otherwise evidence portability becomes an unfunded promise that each supplier and internal function expects someone else to fulfil.

Test the chain before extending delegation

Use the first month of a proposed ninety-day programme to map two consequential workflows. Bring together business owners, legal specialists, operators and complaint handlers. Identify applicable obligations, existing records and missing links. The deliverable is an owned evidence map; the schedule is a management proposal, not a regulatory deadline.

In the second month, test the chain using an authorized historical sample or controlled deployment. Examine explanations, authority limits and recovery. If supplier evidence is unavailable, decide whether to narrow the use or introduce a justified compensating control. Record who accepts the remaining exposure and when that acceptance ends.

In the third month, ask an independent internal reviewer to reconstruct the sample without coaching from its builders. Record retrieval time, missing links, unresolved exceptions and the ability to act on a challenge. Set targets from that baseline. Further delegation should depend on what the organization can explain, constrain and correct.

FROM INSIGHT TO ACTION

Decisions to take forward

LEGAL OWNER AND BUSINESS SPONSOR

Confirm applicable duties before selecting evidence requirements.

Evidence to require: Jurisdiction, activity, role, operative instrument and transition analysis

INDEPENDENT ASSURANCE LEAD

Repair missing links rather than relying on an average compliance score.

Evidence to require: Consequential decisions with actual inputs, authority, reasons and remedy

OPERATING AND PROCUREMENT OWNERS

Approve the boundary only where the claimed response can be exercised.

Evidence to require: Override, correction and portable evidence under supplier change

mayden connects senior judgement, research and engineering to turn strategic choices into working capabilities.

EVIDENCE YOU CAN INSPECT

Research notes

Method and boundaries

Structured comparison of eight opened primary guidance, legal and regulator sources, followed by original operational synthesis; research cutoff 10 October 2026.

Sources were selected for accountability, jurisdiction, assurance and secure-operation relevance. The regulatory examples are explicitly scoped to the EU, US credit and UK data protection.

No regulatory interviews, legal opinion, organization-level audit or causal estimate. Guidance and law have different status. Dates and applicability require confirmation for each deployment; the proposed evidence chain is unvalidated mayden analysis.

Primary sources and author research

[1] AI Risk Management Framework

NIST / 2023-01-26; current page accessed 2026-10-10. Overview of the AI RMF.

[Read the source](#)

[2] Navigating the AI Act

European Commission / Current official FAQ accessed 2026-10-10. Scope and objectives; high-risk systems; application timeline.

[Read the source](#)

[3] When does enforcement start?

European Commission AI Act Service Desk / Current official FAQ accessed 2026-10-10. Transparency and high-risk application dates.

[Read the source](#)

[4] Regulation B, §1002.9 Notifications

Consumer Financial Protection Bureau / Current regulation accessed 2026-10-10. Paragraph (b)(2); official interpretation 9(b)(2)-2 through -7.

[Read the source](#)

[5] Artificial intelligence guidance

Information Commissioner's Office / Current guidance hub accessed 2026-10-10. AI and data protection; explaining decisions.

[Read the source](#)

[6] Introduction to AI assurance

UK Department for Science, Innovation and Technology / 2024-02-12. Sections 3-5, assurance toolkit and practice.

[Read the source](#)

[7] Guidelines for secure AI system development

UK National Cyber Security Centre and international partners / 2023-11-27; version 1.0. Secure design, development, deployment, operation and maintenance.

[Read the source](#)

[8] AI Act, Article 26: Obligations of deployers of high-risk AI systems

European Commission AI Act Service Desk / Consolidated text as at 2026-07-27; accessed 2026-10-10. Article 26(2), (5) and (6). High-risk scope and legal qualifications matter; six months is not a universal retention rule and a control exercise is not legal certification.

[Read the source](#)