Nick Clark · Published April 25, 2026 · PDF

How to Map a Full Governance Chain to GDPR and EU AI Act Obligations

If you run an autonomous system that senses, decides, and acts on the physical world, you eventually have to show a regulator where each obligation is enforced. This guide describes an architectural approach for doing that: a composed governance chain whose properties line up, one at a time, with the obligation categories in GDPR and the EU AI Act. It describes an architecture disclosed in U.S. Provisional Application No. 64/049,409, not a shipping library. The home inventive step is the Five-Property Governance Chain inventive step.

What You Are Building

You are building the compliance-mapping layer for an autonomous system: self-driving vehicles, warehouse robots, port or airfield vehicles, or wearable-equipped personnel that sense their surroundings, coordinate with peers, and actuate in the physical world. Regulators do not ask whether your system "has governance." They ask a narrower, harder question: for each obligation, point to the mechanism that enforces it, and show me the record that proves it fired.

Under GDPR that means naming where lawful basis, purpose limitation, data-subject rights, and accountability live. Under the EU AI Act it means naming where record-keeping, human oversight, transparency, and risk controls live for a high-risk system.

The problem this guide addresses is structural: most autonomous stacks scatter these enforcement points across unrelated subsystems, so a regulator's "where is it enforced" has no single answer.

The approach here is to give every governed action the same five-property spine, then map obligation categories onto those five properties. The properties come from the architecture disclosed in U.S. Provisional Application No. 64/049,409. You implement the mapping yourself; the guide teaches the shape.

Why the Obvious Approaches Fall Short

The common way to reach for compliance is to bolt logging and access control onto an existing sensor-actuator pipeline. This is a real and legitimate practice, and for many systems it is enough. But for an autonomous system that fuses observations from many sources and then moves physical mass, three structural gaps tend to open up.

First, homogeneous authentication. Existing V2X security stacks authenticate messages through public-key infrastructure (PKI) and security credential management systems (SCMS). These are accurately good at proving a message is signed. The filed disclosure notes that they nonetheless treat all authenticated messages homogeneously: a signed message from a regulatory device and a signed message from an unknown peer carry the same weight, because the stack has no authority-taxonomy semantics differentiating behavioral response by the source's governance authority. Lawful-basis and accountability obligations need that differentiation.

Second, provenance as afterthought. In conventional sensor-fusion and data-pipeline output formats, lineage is an optional metadata annotation attached after the fact. The spec contrasts this with a lineage record that is intrinsic to every derived observation, cryptographically attested, and traversable in both directions across any number of intermediate agents. Record-keeping and auditability obligations are far weaker when the record is optional and single-agent.

Third, no admissibility gate before action. Conventional architectures, the disclosure observes (PLCs, SCADA, ECUs, cloud IoT, digital twins), act on anonymous or homogeneously-authenticated signals without a composite admissibility evaluation before actuation. Human-oversight and risk-control obligations have nowhere structural to attach when the sense-to-act path has no gate.

None of this makes the obvious approaches wrong. It means that mapping obligations onto them is a per-subsystem scavenger hunt rather than a lookup.

The Architecture

The disclosed architecture imposes a five-property governance chain on every governed mutation in the system. Per the filed spec, the five properties are:

1. **Authority-credentialed observation.** Each observation carries a credentialed source identification, evaluated through an authority taxonomy. The taxonomy includes levels such as a regulatory-infrastructure authority assigned to specified devices. Source identity here is established through a trust-slope continuity mechanism: identity is validated by continuity of the emitting device's signal rather than by enrollment of a long-lived stored secret, and a continuity-validation output is consumed downstream. This is the identity property.
2. **Evidential weighting in a shared governed observation store.** Observations are weighted by authority, sensing-modality reliability, and inter-source consistency through a composite weighting function. Two devices asserting the same thing do not automatically win; weight is a function of credentialed authority and corroboration.
3. **Composite admissibility evaluation across cognitive domain fields.** Every mutation is evaluated against dispositional, integrity, confidence, and capability fields before admission. The evaluator produces graduated outcomes, not a binary

pass/fail: outcomes disclosed include admit, gate, defer, solicit, reject, handoff, escalate, quarantine, and degraded-mode. This is the policy/admissibility property, and it is the natural attachment point for oversight and risk controls.

4. **Governed actuator execution.** Every physical actuation requires composite admissibility approval before it fires. There is no path from observation to physical movement that skips the gate.
5. **Lineage-recorded provenance.** Every observation, evaluation, and action is linked through deterministic lineage across the architecture. Per the spec the lineage record permits deterministic reconstruction of an observation's provenance back to originating sensor observations, across any number of derivation agents, with every transformation and the governance-policy version recorded, and it is cryptographically attested for tamper-evidence.

Two further properties of the chain matter for a compliance mapping. It is **recursive**: outputs of any primitive re-enter the chain as new authority-credentialed observations, and every actuation passes back through every primitive's governance, so there is no unroverned side channel. And the disclosure includes a **temporal reconstruction** mechanism that rebuilds the view state, filters, and governance policies in force at a prior time, itself producing a governed, lineage-linked observation the spec describes as admissible in regulatory, legal, and governance-enforcement proceedings.

The disclosure also specifies consent and privacy governance as first-class: biometric and operator observations are handled subject to privacy governance, and consent-governance evaluation confirms an action is permitted under applicable privacy and consent constraints before admission. Regulated-domain deployment is called out explicitly, with the derivation-lineage record positioned as the structural element satisfying the requirement that every derived decision be auditable back to its source evidence.

That is the substance you map against. Nothing above asserts a benchmark or a compliance certification; these are the mechanisms the filing discloses.

How to Approach the Build

Work obligation-first, property-first. The goal is a mapping table you can hand a regulator, backed by mechanisms you have actually implemented.

Step 1: Enumerate obligations as testable statements. Write each GDPR and EU AI Act obligation you are subject to as a sentence of the form "the system must be able to show X." Group them into the categories that align with the chain: source/identity, purpose and consent, decision gating and oversight, action control, and record-keeping. Do not skip obligations that have no obvious home; those are where the mapping earns its keep.

Step 2: Bind identity to the authority taxonomy. Implement authority-credentialed observation with a taxonomy that names your real authority levels, including a regulatory-infrastructure level. Use trust-slope continuity for source validation. Map lawful-basis and accountability obligations here, because "who asserted this, under what authority" is now a first-class field rather than a homogeneous signature.

An illustrative interface sketch (not shipping code, faithful to the disclosed fields):

```
# Illustrative only. You implement these.
observation = {
  authority_credential,      # source identity, evaluated via authority taxo
  continuity_validation,     # trust-slope continuity output, not a stored s
  evidential_weight,         # authority x modality-reliability x consistenc
  admissibility_outcome,     # admit|gate|defer|...|reject
  lineage,                   # deterministic, attested, bidirectional
}
```

Step 3: Route purpose and consent through admissibility. Add consent-governance evaluation to the composite admissibility evaluator so that purpose limitation and consent obligations are enforced at the gate, before an observation is admitted or an action approved. Privacy filtering applied at emission time should be recorded so it can be reconstructed later for consent-compliance audit.

Step 4: Make the admissibility evaluator your oversight surface. The EU AI Act's human-oversight and risk-control obligations map onto the graduated outcomes. Configure the evaluator so that the actions a regulator cares about resolve to gate, defer, solicit, escalate, or handoff rather than silent admit, and record the field values (dispositional, integrity, confidence, capability) that produced the outcome.

Step 5: Enforce the actuation gate and prove it. Ensure no actuation path bypasses composite admissibility approval. Then, for each obligation about controlling what the system does, point to the gate and to the lineage entry the gate wrote.

Step 6: Turn lineage into your record-keeping answer. Wire deterministic, attested lineage into every observation and decision, and stand up the temporal-reconstruction query so that "show the state and policy in force at time T" is a single operation. This is your record-keeping, auditability, and litigation-support answer in one mechanism.

Step 7: Produce the mapping table and test it. For every obligation from Step 1, fill in the property, the mechanism, and a concrete query that produces the proving record. An obligation with no query is not yet mapped. Exercise each query end to end against real recorded runs.

The tradeoff to accept going in: this chain touches every governed action, so it is intrusive by design. That intrusiveness is what makes the mapping a lookup instead of a hunt, but it is real engineering cost, and the recursion means you cannot leave an ungoverned fast path for "trusted" internal sources.

What This Does Not Give You

This is an architecture, not a drop-in library. There is no package to install and no SDK behind this guide; you build every mechanism yourself. The pseudocode above is illustrative shape, not a working implementation, and it will not "just work."

It is not a legal compliance determination. The disclosed mechanisms give you enforcement points and provable records; whether a given deployment satisfies a given GDPR article or EU AI Act requirement is a legal judgment for your counsel and, ultimately, a regulator. The architecture makes the mapping expressible and auditable. It does not certify it.

Nothing here is benchmarked or productized. The filing discloses mechanisms and their structure; it does not report performance numbers, and none are claimed here. The consent, privacy-governance, and admissibility rules are governance-policy-configurable, which means their correctness for your jurisdiction is your responsibility to define and test.

Finally, the fit is narrow. This approach targets autonomous, physical-world systems that fuse multi-source observations and actuate. A system with no actuation gate to protect, or with a single trusted data source, gets less from the chain and may be better served by simpler logging and access control.

Disclosure Scope

The governance-chain architecture described here, including the five-property chain, trust-slope continuity identity, composite admissibility evaluation, governed actuation, deterministic lineage, and temporal reconstruction, is disclosed in U.S. Provisional Application No. 64/049,409. This guide is educational: it explains how to approach building and mapping such an architecture onto regulatory obligation categories. It is

not a warranty, not a compliance certification, and not an offer of software. Any implementation, and any determination that it satisfies GDPR or the EU AI Act, is the reader's own responsibility.

Five-Property Governance Chain (/governance-chain)

All 40 steps → (/inventive-steps)

The umbrella primitive: every mutation passes through the same five-property structural test.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- Five-Property Governance Chain: The Architectural Umbrella (/articles/five-property-governance-chain-the-architectural-umbrella)

SECONDARY TECHNICAL

- Authority-Credentialed Observations (/articles/governance-chain/authority-credentialed-observation)
- Evidential Weighting in Governance Chain (/articles/governance-chain/evidential-weighting)
- Composite Admissibility Evaluation (/articles/governance-chain/composite-admissibility)
- Governed Actuator Execution (/articles/governance-chain/governed-actuator-execution)
- Lineage-Recorded Provenance (/articles/governance-chain/lineage-recorded-provenance)
- Recursive Closure Across Governance Chain (/articles/governance-chain/recursive-closure)
- Hierarchical Governance Composition (/articles/governance-chain/hierarchical-composition)
- Technology-Neutral Governance (/articles/governance-chain/technology-neutrality)
- Structural Distinction Test for the Five-Property Governance Chain (/articles/governance-chain/structural-infringement-test)

APPLICATIONS · GENERAL

- Cross-Domain Governance: One Auditable Authority Chain Across Defense, Cyber, Health, and Finance (/articles/governance-chain/cross-domain-governance-umbrella)

- Multi-Jurisdiction Compliance Without a Single Supervening Authority: A Federated Governance Chain (/articles/governance-chain/federated-governance-umbrella)
- Multi-Jurisdictional Compliance for Cross-Border Data Pipelines: A Governance Chain Umbrella (/articles/governance-chain/multi-jurisdictional-governance-umbrella)
- AI Governance Umbrella Across Regulatory Regimes (/articles/governance-chain/ai-governance-umbrella)
- Climate Governance Umbrella (/articles/governance-chain/climate-governance-umbrella)
- Multi-Tier Supply Chain Provenance: One Governance Substrate for NIST 800-161, CISA SSDF, NDAA 5949, CSDDD, and DSCSA (/articles/governance-chain/supply-chain-governance-umbrella)
- A CCPA and CPRA Compliance Architecture for Verifiable Consumer Rights and ADMT (/articles/governance-chain/ccpa-cpra-privacy)
- EU CSDDD Supply-Chain Due Diligence: A Credentialed Governance-Chain Architecture (/articles/governance-chain/eu-csddd-due-diligence)
- CSRD Audit-Ready Sustainability Reporting: A Governance-Chain Architecture for ESRS Assurance (/articles/governance-chain/eu-csrd-sustainability)
- EU Data Act Compliance for Connected-Product Data: A Credentialed Data-Flow Architecture (/articles/governance-chain/eu-data-act)
- NIS2 Compliance Architecture: How to Meet EU 24-Hour Cyber Incident Reporting (/articles/governance-chain/eu-nis2-cyber)
- FedRAMP High and DoD IL5/IL6: Continuous Compliance Evidence as a System Property (/articles/governance-chain/fedramp-il5-il6)
- A GDPR Article 22 Compliance Architecture for Automated Decision-Making and Profiling (/articles/governance-chain/gdpr-article-22)
- HIPAA Security Rule Compliance for Cross-Organization ePHI Access (/articles/governance-chain/hipaa-security-rule)
- How to Produce Auditable IEEE 7000 Series Conformance Evidence (/articles/governance-chain/ieee-7000-series)
- How to Build an ISO/IEC 42001 AI Management System on a Governance-Chain Substrate (/articles/governance-chain/iso-iec-42001-ai-management)
- ITAR and EAR Export Control Compliance: Per-Access Provenance for Deemed Exports (/articles/governance-chain/itar-ear-export-controls)
- NDAA Section 1709 Compliance Architecture: Runtime China-Origin Controls for DoD Supply Chains (/articles/governance-chain/ndaa-1709-china-controls)
- NIST SP 800-53 Rev 5 Compliance for Autonomous and AI-Mediated Systems (/articles/governance-chain/nist-800-53-controls)

APPLICATIONS · SPECIFIC

- AWS IAM Cross-Account vs a Governed Cross-Authority Chain (/articles/governance-chain/aws-iam-cross-account)
- Hyperledger Fabric vs Governed Cross-Authority Composition (/articles/governance-chain/hyperledger-fabric)
- Microsoft Entra ID vs Governed Agent Execution: The Governance Chain Alternative (/articles/governance-chain/microsoft-entra-id)
- AWS Verified Permissions vs a Governed Physical-World Actuation Chain (/articles/governance-chain/aws-verified-permissions)
- CyberArk PAM vs Governed Actuator Execution: The Governance-Chain Substrate (/articles/governance-chain/cyberark-pam)
- Okta Alternative: Governed Cross-Authority Identity Beyond a Single Broker (/articles/governance-chain/okta-identity)
- Ping Identity vs Governed Actuation: A Cross-Vendor Governance Chain (/articles/governance-chain/ping-identity)
- SailPoint IGA vs the Governance Chain: Credentialed Identity Mutation (/articles/governance-chain/sailpoint-iga)
- EU eIDAS 2 and the EUDI Wallet vs a Governed Agent-Execution Chain (/articles/governance-chain/eu-eidas-2-eudi-wallet)
- What happens after Microsoft Entra Verified ID verifies a credential? (/articles/governance-chain/microsoft-entra-verified)
- Pollen Mobile vs Governed Coverage Evidence: A DePIN Governance Chain (/articles/governance-chain/pollen-mobile)

Five-Property Governance Chain overview → (/governance-chain)