AIVR
[Sovereign AI Platform]

THE COMPLIANCE BRIEF · v1.0

Compliant by Construction.

Eight simultaneous EU regulatory regimes.
One architecture. Zero retrofits.

AI Act · GDPR · Cyber Resilience Act · Digital Services Act
Data Act · Data Governance Act · NIS2 · Product Liability Directive

PREPARED BY	FOR	DATE
AIVR Architecture Team	**Legal, Procurement, IT Security**	**11 April 2026**

01 THE URGENCY

August 2 Is Basically Tomorrow

2 August 2026 is when the EU AI Act's high-risk obligations come into force across the Union. The general-purpose AI model rules already kicked in on 2 August 2025. Most GenAI vendors are still retrofitting compliance. **AIVR was built for this moment from version 1.**

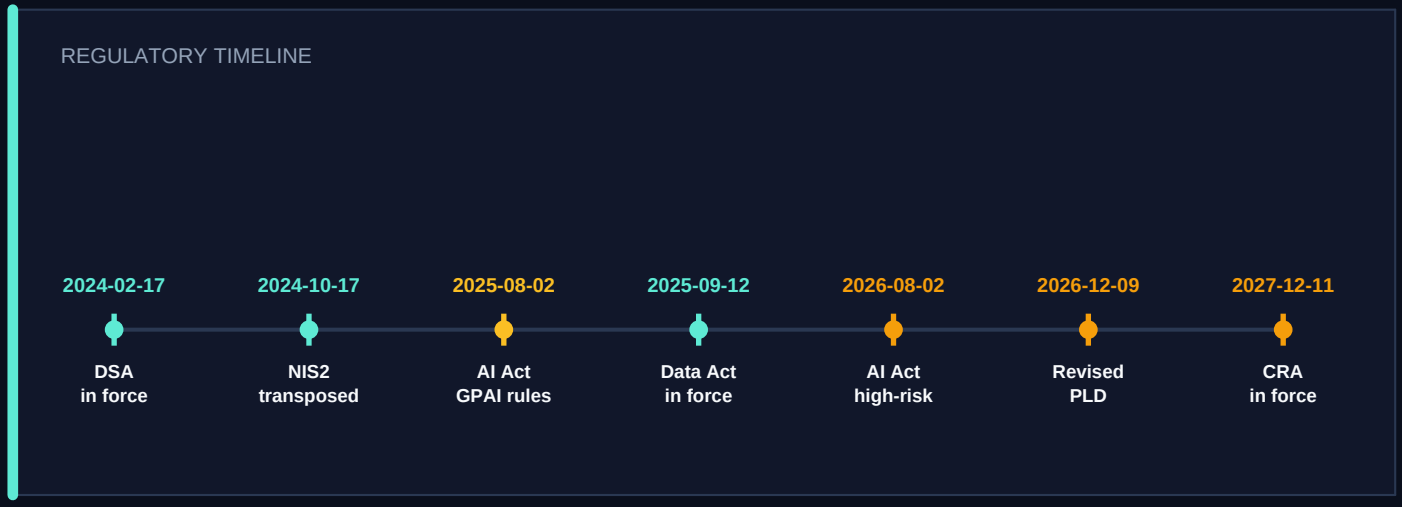

02 THE POSITION

Compliance Is an Architecture, Not a Checklist

Most vendors approach EU compliance as paperwork bolted on after the product is built — privacy notices, DPIAs, and acceptable-use policies layered on top of a system that was never designed to honour them. The model breaks the moment a regulator asks *"show me the audit trail"* or *"prove the human was in the loop."*

AIVR inverts the order. Every compliance obligation we tracked maps to a concrete component already in the architecture. The Sentinel Guardian is human-in-the-loop. The Write-Ahead Log is post-market monitoring. The Replay Engine is evidence disclosure. The per-project isolation is data minimisation. **None of these were added for compliance. They were the architecture.** Compliance is what falls out the other end.

03 FIVE COMPLIANCE PILLARS

What AIVR Actually Delivers

Every line of the eight-regime table on the next page resolves to one of these five pillars. Each pillar is a concrete component in the v3.4.0 architecture, running in production today.

01

Sovereignty by default

Your data, your models, your iron. No third-party processors, no cross-border transfers, no shadow telemetry. Eliminates the GDPR Chapter V transfer problem and the Data Act vendor lock-in trap in one architectural decision.

02

Auditability by construction

Every agent action is captured in a per-agent Write-Ahead Log and the NATS REPLAY_LOG stream (30-day retention, indexed by agent, task, and timestamp). The Replay Engine reconstructs any decision step-by-step.

03

Human oversight, built in

The Sentinel Guardian intercepts every risky action via NATS and gates it against ADRs and prerequisites. SLB enforces dual-key approval on dangerous commands. WebUI Permission Handler surfaces every tool call. AI Act §14 satisfied by construction.

04

Risk classification automated

DCG blocks 45+ destructive patterns. SLB applies a four-tier risk model (SAFE / CAUTION / DANGEROUS / CRITICAL). The Sentinel runs continuous safety intent analysis on dedicated NPU silicon. AI Act §9 implemented in code, not policy documents.

05

Resilience and recovery

The Sentinel is a native Windows service that survives Docker, NATS, and network outages. Its skill catalog contains 24 recovery runbooks. NIS2 Article 21 cybersecurity risk management is not a policy — it is an executable.

04 THE EIGHT REGIMES · PART I

What Each Regulation Demands (1 of 2)

Every row below is a regulation in force or arriving inside the next 18 months. Every row maps to concrete AIVR components by section reference. Nothing is aspirational — these are architectural pillars in the v3.4.0 spec, in production today. Part II continues on the next page with the remaining four regulations.

Regulation	Effective	Demands	How AIVR Delivers
EU AI Act (2024/1689)	**GPAI 2 Aug 2025** **High-risk 2 Aug 2026**	Risk classification, transparency, human oversight, post-market monitoring, incident reporting, technical documentation, automatic logging.	**S47.5** Sentinel Guardian = human-in-the-loop gates. **S43** WAL + Replay Engine = post-market monitoring + automatic logging. **S37** AgentMetrics = technical documentation. **S42** RCA Loop = incident reporting. **S11** SLB four-tier = risk classification.
GDPR (2016/679)	**In force 25 May 2018**	Data minimisation, purpose limitation, right to erasure, data subject access, lawful basis, DPIA, 72h breach notification.	**S13** per-project isolation = data minimisation. **S44** 3-tier archive = right to erasure. NATS audit log = data subject access trail. Sovereign hosting eliminates third-party processor risk. **S47.9** Sentinel incident reports handle 72h notification.
Cyber Resilience Act (2024/2847)	**11 Dec 2027**	Vulnerability handling, security-by-design, SBOM, 24h incident reporting, secure-by-default, free security updates.	**S11** DCG + SLB = secure-by-default. **S11.4** per-agent sandboxing = security-by-design. **S47.6** Sentinel catastrophic recovery = vulnerability handling. AI Vault SHA-256 = SBOM-grade supply chain integrity.
Digital Services Act (2022/2065)	**17 Feb 2024**	Content moderation transparency, illegal content removal, ad transparency, recommender system transparency.	Sentinel content moderation layer. Audit log captures every removal. **S16.3** reward signals document why models routed where they did. **S43** Replay Engine provides full algorithmic transparency.

04 THE EIGHT REGIMES · PART II

What Each Regulation Demands (2 of 2)

Part II of the regime map covers the Data Act, Data Governance Act, NIS2, and the revised Product Liability Directive. The liability window on the revised PLD is ten years, which makes AIVR's built-in evidence disclosure via the Replay Engine a meaningful durable advantage.

Regulation	Effective	Demands	How AIVR Delivers
Data Act (2023/2854)	**12 Sep 2025**	Data sharing fairness, switching/portability, IoT data access, B2G data sharing, contractual fairness for SMEs.	**S20** module-based architecture = open formats (NATS, Qdrant, Postgres, JSONL). Export at any time. No vendor lock-in by construction. AI Vault rollback supports clean switching.
Data Governance Act (2022/868)	**24 Sep 2023**	Data intermediation services, data altruism, consent management, technical and legal interoperability, neutral intermediation.	Sovereign hosting model. **S13** Project Segregation = legal interoperability. Consent state in Redis with full audit trail. The Cockpit is a neutral intermediation layer.
NIS2 (2022/2555)	**Transposed 17 Oct 2024**	Cybersecurity risk management, 24h/72h/1-month incident reporting, supply chain security, governance accountability, training.	**S47.6** Sentinel health monitoring + recovery. **S42** RCA Loop auto-generates incident reports. Supply chain via AI Vault SHA-256. **S47.9** Sentinel External Coordination handles 24h notification.
Revised Product Liability Directive (2024/2853)	**9 Dec 2026**	Strict liability for AI/software, defect presumption when evidence is withheld, mandatory evidence disclosure, 10-year liability window.	**S43** Replay Engine = evidence disclosure on demand. WAL = audit trail of every decision. **S16.3** reward scoring documents "reasonable care." **S47.5** Sentinel Guardian gates = defect prevention upstream.

05 WHAT LEGAL ACTUALLY ASKS FOR

The Five Questions, Answered Up Front

When a regulator, auditor, or in-house counsel sits down with an AI vendor, the conversation always converges on the same five questions. Here is how AIVR answers each one without leaving the architecture document.

Q1 — Can you reproduce exactly what the agent did, when, and why?

Yes. Every action is appended to a per-agent Write-Ahead Log and mirrored to the NATS REPLAY_LOG stream (100K buffer, 30-day retention). The Replay Engine has five modes: full, range, step-through debugger, failure-focused, and diff replay. Satisfies AI Act §12, NIS2 incident evidence, and revised PLD evidence disclosure — all from the same component.

Q2 — Where is the human in the loop?

Five places. (1) WebUI Permission Handler surfaces every tool call. (2) SLB classifies every command SAFE/CAUTION/DANGEROUS/CRITICAL with dual-key auth on the top tier. (3) Sentinel Guardian gates risky actions against ADRs. (4) Blueprint approval requires human sign-off. (5) Self-Optimization requires human approval before applying any policy update. AI Act §14 — done.

Q3 — How do you prevent and report incidents?

The Sentinel runs catastrophic recovery as a native Windows service that survives Docker, NATS, and network failures. Its skill catalog contains 24 runbooks. The RCA Loop classifies every novel failure with Gemini Ultra and emits a structured incident report in minutes. Sentinel External Coordination feeds NOC/DCIM via webhook within the NIS2 24-hour early-warning window.

Q4 — Where is the data, and who can touch it?

On your iron. Period. The Cockpit, NATS, Qdrant, Redis, Ollama, and the Sentinel all run on hardware you own. Per-project isolation prevents cross-contamination at the storage layer. Per-agent sandboxing prevents lateral movement. AI Switch routes outbound calls through AI Meter for full token-level accounting. Nothing leaves without appearing in your audit log first.

Q5 — Can you produce evidence on demand?

Yes, in three formats. (1) The Replay Engine outputs full execution traces as JSONL. (2) AgentMetrics exports 12 KPIs as Prometheus time-series with 90-day retention. (3) Compliance Reports generates SOC2/GDPR/incident PDFs on a schedule and on demand. All evidence is signed with project-scoped keys and stored on hardware you control.

06 THE BOTTOM LINE

Most GenAI vendors will spend 2025 through 2027 retrofitting compliance onto architectures that were never designed to honour it.

AIVR was built for this moment from version 1.

8	47	0	100%
REGIMES	SECTIONS	RETROFITS	SOVEREIGN
Mapped to one architecture	v3.4.0 architecture references cited	Compliance is the design, not a bolt-on	Your iron. Your data. Your audit trail.

Sovereignty, auditability, human oversight, risk classification, and resilience are not features we shipped after a regulator complained. They are the architecture. The eight regimes above did not change what we build; they changed who else has to catch up. If your legal team is staring at a 2026-08-02 deadline and a vendor list full of question marks, this is the document to put in front of them.

REFERENCES

arxiv.org/abs/2604.04604 — the eight-regime interaction analysis cited throughout this brief. Maps every inter-regulation overlap and gap. Worth reading in full.

AIVR Technical Architecture v3.4.0 — the 11-part split master document, 47 sections, 21-phase engineering roadmap. All section references in this brief resolve here.

AIVR Cockpit UI Specification v1.0.0 — the operator interface that surfaces every audit, every decision, every reward signal cited above.