TAUDIL

Deterministic Intelligence Layer

What It Is • What Pain It Resolves • What It Can Do

Version 1.0 | May 2026

1. What Is TauDIL

TauDIL — Deterministic Intelligence Layer — is a governance platform that sits between your existing systems and the AI models they use. It ensures that every AI-assisted decision in your organisation is safe, explainable, compliant, and human-controlled — before it reaches anyone who acts on it.

TauDIL does not replace your AI models. It does not replace your business systems. It governs them. It wraps around every AI interaction — every prompt, every output, every decision — and applies a deterministic set of rules, checks, and safeguards that you define. The result is traceable to a specific rule or policy, explainable in plain language, and logged immutably for audit at any future point.

> TauDIL is model-agnostic. It works with Anthropic Claude, Mistral, Cohere, Gemini, Intellect-3, and any LLM accessible via API. The governance layer is always deterministic regardless of which AI model is underneath.

The name reflects the design philosophy. Tau (τ) is the mathematical constant of full rotation — completeness. DIL — Deterministic Intelligence Layer — reflects that the governance logic is deterministic. Same input always produces the same governance outcome. There is no probability in the safety net.

1.1 The One-Sentence Version

> **TauDIL is the governance layer that ensures AI in regulated industries does what it is supposed to do, nothing more, and proves it — every time.**

1.2 Two Pillars

TauDIL operates across two pillars that work together as a single governance system:

Pillar 1

AI Governance

Controls the AI model itself — what it can say, how coherent it is, whether it is drifting from its purpose, and what it has done before.

- → TRCP-Φκ coherence scoring
- → TRVC triple-redundancy consensus vote
- → DriftAligner session alignment
- → USE 6-layer prompt security
- → Aelthered Mirror per-user AI ledger

Pillar 2

Business Governance

Controls business decisions — who can approve what, when a human must review, what the organisation knows, and what the regulator can see.

- → UAE deterministic assessment engine
- → Authority matrix + escalation with timeouts
- → CKG-RAG domain knowledge graph
- → 44-control compliance engine
- → Aelthered Chronicles immutable audit trail

2. The Problem TauDIL Solves

Organisations are deploying AI at scale in regulated environments without adequate governance. The consequences are real: wrong decisions, unexplainable outputs, regulatory exposure, and zero audit trail when something goes wrong.

This is not a problem of bad AI models. The models are increasingly capable. The problem is the absence of a governance layer between the model and the decision. When an AI model produces a credit recommendation, an underwriting assessment, a candidate screening result, or a clinical suggestion — there is typically no deterministic check that the output is coherent, that it complies with policy, that a human reviewed it, or that anyone can explain it six months later when the regulator asks.

2.1 The Seven Pains

Pain 1 — Hallucination in high-stakes decisions

AI models fabricate facts, cite non-existent regulations, invent risk factors, and produce confident wrong answers. In banking, insurance, healthcare, and law — a hallucinated fact in a decision is a liability. There is no way to detect it at the point of decision without a coherence scoring layer. By the time it is discovered, the decision has been made.

Pain 2 — No explainability when regulators ask

EU AI Act Article 13 requires transparency. GDPR Article 22 requires explanation of automated decisions. MiFID II requires best execution evidence. Basel III requires audit trails. When a regulator asks "why did your AI make this decision on this date?" — most organisations cannot answer. The AI model has no memory. The logs do not capture the reasoning. The decision is undefendable.

Pain 3 — Humans bypassed or overwhelmed

AI systems are supposed to support human decision-making, not replace it. In practice, when AI outputs arrive with high confidence scores and no friction, humans rubber-stamp them. When AI flags too many false positives, humans stop reviewing. Neither outcome is acceptable in regulated environments. There is no structured escalation path that routes the right decisions to the right humans with the right information and a deadline.

Pain 4 — AI systems manipulated through their inputs

Prompt injection, jailbreaking, goal drift, and adversarial inputs are not theoretical. They are observed in production systems. A malicious actor can manipulate an AI underwriting assistant through a client-submitted form. A claims processor can drift from policy lookup into settlement advice. An HR AI can be manipulated to bypass screening criteria. Without a security layer on AI inputs and outputs, the AI becomes a vulnerability.

Pain 5 — No domain-specific knowledge — AI answers from training data, not your policies

An AI model trained on public data does not know your internal underwriting authority matrix. It does not know your organisation's approved risk thresholds. It does not know your regulatory obligations specific to your jurisdiction and licence. Without a domain knowledge layer, the AI is answering from generic training data — not from what your organisation has decided and documented.

Pain 6 — Compliance is a spreadsheet, not a system

Most organisations track compliance manually — a spreadsheet of controls, updated quarterly by a compliance officer. When a regulator asks for evidence, someone extracts it by hand. When a control fails, it is discovered in the next quarterly review, not in real time. There is no continuous, automated, evidence-backed compliance monitoring tied to live system state.

Pain 7 — EU AI Act deadline with no implementation plan

The EU AI Act applies from August 2026 for high-risk AI systems. Organisations in financial services, insurance, HR, healthcare, and government services are operating high-risk AI systems today. The Act requires a risk management system, technical documentation, record-keeping, human oversight controls, accuracy and robustness measures, and a Fundamental Rights Impact Assessment. Most organisations have none of these in place as implemented systems — only policies on paper.

3. How TauDIL Resolves It

TauDIL addresses each of these pains directly and specifically. Not with a policy framework. Not with guidance documents. With implemented, running, deterministic system components.

✓ Pain 1 resolved — TRCP-Φκ coherence scoring blocks hallucinated outputs

Every AI response is scored across three independent lenses: semantic entropy (how uncertain the response is), contradiction detection (whether it contradicts verified knowledge in the CKG), and goal overlap (whether it aligns with the session's declared purpose). Responses below the coherence threshold are blocked before they reach the user. The kappa score is logged with every decision.

✓ Pain 2 resolved — Aelthered Chronicles: every decision is explainable and verifiable

Every governance event — every rule that fired, every escalation, every AI output hash, every compliance assessment — is logged to Aelthered Chronicles. Each record is ED25519-signed and SHA-256 hash-chained. The chain can be verified independently by a regulator without involving TauGuard. When a regulator asks "why was this decision made on this date?" — the answer is in the chain, immutable, with a cryptographic proof of integrity.

✓ Pain 3 resolved — UAE escalation engine with authority matrix and timeouts

The UAE (Unified Assessment Engine) routes every REVIEW and BLOCK verdict to the correct human in the authority matrix — by domain, by decision type, by authority level. Each escalation has a countdown timer. If the assigned reviewer does not act within the SLA, the escalation auto-escalates to the next level. No AI verdict sits unreviewed. No human is bypassed. The escalation chain is configurable by domain owner and auditable.

✓ Pain 4 resolved — USE 6-layer prompt security blocks manipulation at the input and output

The Unified Security Engine evaluates every AI interaction across six layers: prompt injection detection, jailbreak pattern matching, goal drift detection, data exfiltration prevention, PII leakage detection, and cross-domain data boundary enforcement. Blocked interactions are logged to Aelthered Mirror — the per-user AI behaviour ledger — so patterns of manipulation attempts are visible to the security team.

✓ Pain 5 resolved — CKG-RAG: AI answers from your verified organisational knowledge

The Company Knowledge Graph (CKG) stores your organisation's verified entities, policies, regulations, procedures, and relationships in a domain-isolated knowledge graph. When the AI responds to a query in the Insurance Underwriting domain, it draws from your Solvency II documentation, your underwriting authority matrix, and your internal risk thresholds — not from generic training data. The AGL (Admission Gate Layer) activates automatically when the CKG reaches sufficient quality, ensuring the AI only operates within verified knowledge boundaries.

✓ Pain 6 resolved — ComplianceEngine: continuous real-time compliance monitoring against 8+ frameworks

The ComplianceEngine runs every 24 hours and on-demand, assessing live system state against EU AI Act, GDPR, ISO 27001, SOC 2 Type II, Solvency II, Basel III, MiFID II, and HIPAA. Every control check queries real system data — not documentation. When a control fails, it is visible immediately in the domain dashboard with severity rating and specific remediation guidance. The evidence chain is automatically available for export to regulators.

✓ Pain 7 resolved — EU AI Act compliance built in, not bolted on

TauDIL is designed around the EU AI Act. The FRIA gate enforces Article 27 before any high-risk domain activates. The human oversight controls satisfy Article 14. The audit trail satisfies Article 12. The coherence scoring satisfies Article 15. The technical documentation is maintained by Aelthered Chronicles. The system-wide EU AI Act coverage is 94.4% out of the box, with a clear path to 100%.

4. What TauDIL Is Capable Of

TauDIL is a production system. The following capabilities are operational today.

	Capability	What it does
⚖️	**Deterministic assessment**	Evaluates any structured payload against configurable domain rules. Returns APPROVE / REVIEW / BLOCK in <500ms. Same input always produces same output. No probabilistic governance.
🔺	**Escalation with authority matrix**	Routes REVIEW and BLOCK verdicts to the correct human by domain, decision type, and authority level. Countdown timers enforce SLA. Auto-escalation if unreviewed. Full chain logged.
🧠	**AI coherence scoring**	TRCP-Φκ scores every AI output across entropy, contradiction, and goal overlap. Outputs below threshold blocked. Score logged with every decision. Kappa metric visible to domain owner.
🛡️	**Prompt security**	USE 6-layer blocks prompt injection, jailbreaking, goal drift, data exfiltration, PII leakage, and cross-domain boundary violations. Real-time. No AI model sees a harmful input if TauDIL blocks it first.
📚	**Domain knowledge grounding**	CKG-RAG builds a domain-isolated knowledge graph from your documents. AI responses grounded in verified organisational knowledge, not generic training data. AGL activates when quality sufficient.
📋	**Rules as code**	Business and compliance rules encoded in TauDIL DSL. Domain-specific, department-scoped, priority-ordered. AI and business rules separated. All rule changes logged with changelog to Aelthered Chronicles.
🔒	**Immutable audit trail**	Aelthered Chronicles: ED25519-signed, SHA-256 hash-chained record of every governance event. Tamper-evident. Independently verifiable by regulators. Full chain export available.
🪞	**Per-user AI behaviour ledger**	Aelthered Mirror: per-user immutable record of AI outputs, manipulation attempts, goal drift events, kappa drops. AI reads its own history before each session. Patterns of behaviour visible to security.
📊	**Continuous compliance monitoring**	ComplianceEngine assesses live system state against 8+ frameworks every 24 hours. Scores per control. Missing requirements identified with severity and remediation. Evidence exportable for regulators.
🏛️	**FRIA gate**	Fundamental Rights Impact Assessment gate enforced before any high-risk AI domain activates. EU AI Act Article 27 compliance. Results stored with timestamp and signatory. Cannot be bypassed.
🌐	**Multi-domain isolation**	Each business domain (Credit Risk, Insurance, AML, KYC, Trade Finance, Recruitment) is isolated. Different rules, different CKG, different authority matrix, different compliance configuration. Cross-domain access requires explicit rules.
🔗	**REST API + webhooks**	Any system that can make an HTTPS POST can integrate with TauDIL. HMAC-SHA256 signed webhooks push verdicts asynchronously. <500ms synchronous assessment response. On-premises — zero data egress.

🤖	**AI assistant per domain**	Each domain has an AI assistant grounded in the domain CKG. Queries answered from verified organisational knowledge. Every response scored, logged, and subject to governance rules. Not a general-purpose chatbot.
👤	**Auditor Intelligence**	Natural language query over the audit trail. "How many queries were blocked this month?" "What happened in the Insurance domain last week?" "Show me all escalations above Level 5." Deterministic answers from live data.
📦	**On-premises deployment**	Runs on the client's own infrastructure. No data leaves. No cloud dependency. No vendor lock-in for data. Air-gapped deployment possible. Models run locally. Audit chain stays on client server.

5. Where TauDIL Fits — By Industry

TauDIL is built for regulated industries where AI decisions have real-world consequences. The following table maps the primary pain in each sector to the specific TauDIL capability that resolves it.

Sector	Primary Pain	TauDIL Resolution
Banking & Credit	AI credit scores are unexplainable. Basel III audit trail missing. No human oversight on automated decisions.	UAE assessment engine + escalation matrix. Aelthered Chronicles audit trail. TRCP coherence scoring on every recommendation.
Insurance Underwriting	AI underwriting outputs unchecked for coherence. No Solvency II evidence chain. Underwriting authority bypassed.	CKG grounded in Solvency II. Underwriting rules with authority levels. BLOCK verdicts trigger escalation to Principal Underwriter.
AML / Financial Crime	High false positive rates from AI. SAR generation unaudited. Transaction structuring not detected systematically.	Rule engine with FATF-aligned patterns. Deterministic APPROVE/REVIEW/BLOCK. Every SAR trigger logged to Aelthered Chronicles.
KYC / Identity	AI identity verification not auditable. PEP/sanctions screening gaps. GDPR Art.22 automated decision safeguards absent.	Cross-domain rules enforce sanctions checks. GDPR Art.22 escalation path built in. Full audit chain per verification decision.
HR / Recruitment	AI hiring decisions unexamined for bias. EU AI Act Annex III §4(a) high-risk. No FRIA. No authority matrix.	5-phase pipeline with bias-exclusion SafeConditionEvaluator. FRIA gate. Authority matrix: Recruiter → HR Head → CEO.
Healthcare	AI diagnostic outputs presented to clinicians without coherence check. Patient data crosses department boundaries.	TRCP coherence gate before clinical presentation. Strict domain isolation. Authority matrix: nurse → doctor → consultant → CMO.
Legal	AI contract analysis outputs not validated against firm knowledge. Client matter confidentiality not enforced for AI.	CKG stores legal precedents and obligations. TRVC validates AI output against CKG. Cross-domain rules enforce matter confidentiality.
Government	AI eligibility decisions lack human	UAE assessment with mandatory

	review path. GDPR Art.22 violated. No tamper-proof audit trail for public accountability.	human sign-off. Art.22 escalation built in. Aelthered Chronicles provides public accountability chain.

6. How It Integrates — Without Replacing Anything

TauDIL does not require replacing your existing systems. It sits as a governance layer that your existing systems call before committing decisions. The integration is a single HTTPS POST.

> Existing system → POST payload to TauDIL REST endpoint → TauDIL runs governance pipeline → returns APPROVE / REVIEW / BLOCK + evidence → existing system acts on verdict. TauDIL never initiates calls to external systems. All data processed on your server. Zero data egress.

Integration takes hours, not months. Any system that can make an HTTPS POST request — SAP, Workday, Guidewire, Temenos, Epic, Cerner, Bloomberg, custom platforms — can integrate with TauDIL via the REST API or receive verdict events via HMAC-SHA256 signed webhooks.

6.1 What TauDIL Is Not

To be precise about the boundaries:

- **TauDIL is not an AI model.** It governs AI models. It is model-agnostic and works alongside any LLM.
- **TauDIL is not a compliance consultancy.** It is an implemented system. The rules are code. The audit trail is cryptographic. The compliance score is computed from live system state, not from interviews.
- **TauDIL is not a replacement for your core business system.** It wraps around it. Your existing platform continues to operate. TauDIL adds the governance layer on top.
- **TauDIL is not a cloud service (at this stage).** It runs on your infrastructure. Your data stays on your server. This is a deliberate design decision for regulated industries where data sovereignty is non-negotiable.

7. Summary

TauDIL in one page

The governance layer between your organisation and its AI systems.

The problem:

Organisations are deploying AI in regulated environments with no governance layer. AI hallucinations reach decision-makers unchecked. Outputs are unexplainable to regulators. Humans are bypassed. There is no audit trail. Prompt injection and manipulation go undetected. AI answers from training data, not from the organisation's policies. The EU AI Act applies from August 2026.

The solution:

TauDIL sits between your existing systems and your AI models. It scores every AI output for coherence, applies your rules deterministically, routes decisions to the right humans with SLA enforcement, grounds AI responses in your verified organisational knowledge, and logs everything immutably with cryptographic proof. Same input always produces same governance outcome.

Key facts:

→ On-premises — your data stays on your server. Zero egress.
→ Model-agnostic — works with any LLM via API.
→ Deterministic — same input always produces same governance outcome.
→ Auditable — every decision traceable to a specific rule, logged with ED25519 signature.
→ Compliant — 94.4% EU AI Act coverage out of the box.
→ Integrates via REST API — any system that can make an HTTPS POST can connect.
→ Active today — AGL live, 6 domains governed, 1,500+ audit events logged.