Product factsheet

AI Governance Platform

Software solution

Description

RegenData is a constraint enforcement and audit layer for AI systems operating in water and physical infrastructure domains.

AI systems operating in physical infrastructure can produce confident, fluent outputs that violate the physical constraints of the system they are advising. Existing tools, sensors, control systems, document retrieval, digital twins, improve what AI knows. They do not govern what AI is permitted to conclude from that information.

RegenData addresses that gap through three components. A domain canon encodes the physical, chemical, and biological mechanisms governing a specific infrastructure domain as a structured rule set. An execution boundary evaluates every AI output against the active canon before it reaches an operator or system, and either permits it or blocks it and routes it to human review. A provenance store seals an immutable, tamper-evident record of every gate decision, capturing the input, canon version, constraint rules evaluated, and outcome.

The first deployed canon covers onsite wastewater treatment systems. Pilot validation demonstrated a 76% reduction in unsafe AI recommendations under canon-constrained versus unconstrained conditions using the same model and the same prompts. The constraint layer was the only variable.

RegenData is currently seeking co-development partners in water and wastewater domains to build and deploy canons for adjacent physical infrastructure verticals. The EU AI Act classifies water and wastewater AI as high-risk systems. Enforcement begins August 2026.

The execution boundary and provenance store are currently in active development. Co-development pilot engagements are open now

Target audience

Water utilities, wastewater operators, infrastructure AI companies, engineering firms, and risk and compliance teams responsible for deploying or governing AI in water and physical infrastructure systems.

Actors, their roles and interactions

Water utilities and wastewater operators: deploy RegenData as the constraint and audit layer governing AI-assisted operational decisions before they reach control systems or field operators.

Infrastructure AI companies and platform providers: integrate the RegenData execution boundary as a post-inference constraint layer within existing AI deployments, digital twins, and decision support platforms.

Engineering firms and system integrators: embed the canon architecture within AI-assisted design, assessment, and compliance workflows.

Regulators and compliance bodies: use the RegenData provenance store as a tamper-evident audit trail demonstrating that AI outputs were governed against a versioned domain canon at the moment of decision.

Insurers and risk underwriters: use the atomic decision records to establish a clear, reconstructable liability chain for AI-mediated infrastructure decisions.

Unique selling points

- Deterministic constraint layer: AI outputs are permitted or blocked against domain rules, not probabilistically weighted
- Mechanism-anchored: constraint rules are grounded in the physical, chemical, and biological reality of the domain, not content filters
- Atomic provenance records: every gate decision generates an immutable, tamper-evident record capturing input, canon version, constraint rules evaluated, and outcome
- Works across existing stacks: sits between AI output and real-world action regardless of the underlying model, data source, or platform
- Validated result: 76% reduction in unsafe AI recommendations in pilot study, same model, same prompts, constraint layer as the only variable
- Regulator and insurer ready: audit trail satisfies EU AI Act high-risk AI documentation requirements
- No retraining required: deploys at inference, not at training

Technical requirements

RegenData deploys as a post-inference middleware service via API. Integration requires the ability to route AI outputs through the execution boundary before delivery to operators or systems. Canon deployment is a configuration load. No model retraining or modification required.

Software data

- Version: 0.1
- Initial release: 2026
- Operating environments:
  - SaaS - Web application
- License type: Commercial

Publications

Box, C. (2025). Mechanism-Anchored Constraint Architecture (MACA): A Framework for Governing AI Outputs in Physical Infrastructure Domains. RegenData Technical White Paper. https://github.com/Colin-io/RegenData-v2-whitepaper/blob/main/RegenData_V2_Technical_Paper.pdf

URL

https://regendata.ai

Technologies applied by the product

- **AI techniques**
- **Hazard Identification and Characterization**
- **Risk Assessment and Management**

Costs

Co-development pilot engagements are structured on a project basis. Licensing terms for deployed canons are defined based on domain scope and deployment scale. Contact to discuss.

Related tags

Digital Monitoring System | Data-Driven Water Governance | Real-Time Data Monitoring

Downloads

The following file can be downloaded from the online page of the product: https://mp.watereurope.eu/d/product/212

- RegenData MACA Technical White Paper
  Mechanism-Anchored Constraint Architecture framework documentation.