When agentic AI crosses a border

a unified governance architecture for data sovereignty, cloud strategy and AI compliance in energy operations

Conference Category	AI & Digital Transformation
Submission Type	Technical Paper / Oral Presentation
Industry Context	Cloud & Digital Infrastructure for O&G — Cross-Sector
Keywords	Data sovereignty · AI agents · Cloud compliance · Industrial AI · Governance-as-a-Service
Author	Koceila Barchiche (LinkedIN) : Chief Data Officer - Atelic AI

1. Objectives

AI agents are being deployed across cloud-connected energy infrastructure to optimise production, manage logistics & inform capital decisions. These 'agents' often traverse sovereign boundaries, multi-region cloud environments & NOC-regulated data zones with no unified governance layer in place.

Existing frameworks address either data governance for static assets or cloud compliance as an IT workstream; neither is designed for the speed, autonomy and cross-domain context-sharing of agentic AI pipelines.

This paper argues that data sovereignty, cloud architecture and AI compliance are not separate workstreams, rather they are three expressions of a single, unresolved governance problem. Resolving these three independently produces compliance failures that are architectural in nature and costly to remediate after deployment.

The paper sets out to:

- Define the structural dependencies between these three dimensions in integrated energy operations;
- Present a reference three-layer governance architecture applicable to cloud-connected O&G operators;
- Introduce two implementable governance constructs (resource-bounded agent behavioural constraints and infrastructure-layer trust scoring)
- Demonstrate their application through a structured scenario analysis; and
- Identify the regulatory obligations most directly implicated, including the UAE Personal Data Protection Law (PDPL), the EU AI Act, and the NIST AI Risk Management Framework.

2. Methods and Procedures

The governance architecture is constructed through a structured review of emerging AI governance research; including recent multi-agent coordination literature addressing behavioural constraint formalisation. This, combined with scenario analysis grounded in cloud infrastructure patterns is representative of mid-to-large integrated energy operators.

The scenario models a hybrid cloud deployment (multi-region, multi-provider) with agents operating across upstream production optimisation, HSE anomaly detection and supply chain coordination, with data flows crossing sovereign and contractual boundaries. The architecture is organised into three interdependent layers:

Layer	Governance Dimension	Industrial AI Failure Mode
L1 — Data	Data Sovereignty	Residency constraints, cross-border transfer restrictions and regulatory data classification expressed as network topology and IAM policy, not application-layer filters
L2 — Infrastructure	Cloud Strategy	compliance constraints encoded at architecture design rather than applied retroactively
L3 — Agency	AI Compliance	auditable agent behaviour through declared resource bounds including token budgets, iteration caps and execution time limits, and an external trust-scoring layer operating as infrastructure middleware independent of the underlying model.

The trust-scoring layer evaluates agent outputs against a configurable threshold derived from compliance history and violation severity, with graduated interventions from advisory flagging through to blocking downstream propagation. It is presented as illustrative configuration values rather than empirically validated parameters.

The scenario traces data through the full pipeline from raw sensor ingestion to decision output, identifying at each stage which sovereignty constraint applies, which cloud architecture choice is implicated and which compliance mechanism is engaged.

3. Results, Observations, and Conclusions

Scenario analysis produces three principal observations.

1. First, when data sovereignty is applied as a policy overlay after cloud architecture is designed, enforcement becomes dependent on manual controls that fail at the speed and volume at which agents operate; sovereignty constraints must be expressed architecturally before any agent is deployed against that data.
2. Second, agents operating without declared behavioural bounds produce non-deterministic resource consumption and output behaviour that cannot be audited against a compliance baseline; agents with declared constraints produce auditable behaviour profiles, creating a meaningful distinction for compliance and CDO functions.
3. Third, trust scoring implemented as infrastructure middleware, rather than inside the agent model, can enforce compliance thresholds without model modification and can be governed independently by a compliance or risk function.

The paper concludes that operators who treat these three layers as separate procurement or IT workstreams will face compliance exposure that is structural and expensive to address post-deployment, and that the organisations best positioned for agentic AI at scale are those that resolve the governance architecture before agent deployment begins.

---

4. Novel and Additive Information

This paper's primary contribution is the framing of data sovereignty, cloud strategy and industrial AI compliance as a unified design space. This is distinct from existing data governance frameworks, which address static assets, and from cloud compliance literature, which does not account for autonomous agent pipelines.

The three-layer reference architecture provides a replicable model that compliance, architecture and CDO functions can adopt without modification of underlying AI models. The application of resource-bounded behavioural constraints and infrastructure-layer trust scoring to the specific regulatory and operational context of cloud-connected energy infrastructure, including UAE PDPL, EU AI Act and NIST AI RMF obligations; has not been addressed in previous existing ADIPEC technical literature.

The worked scenario provides a structured reference that technical teams can use to assess the compliance readiness of current agentic deployments against concrete architectural criteria.
