emma

Architecture Patterns for Sovereign Workload Placement

How European mid-market companies can govern workloads across hyperscaler and sovereign infrastructure through policy-driven placement, continuous compliance evidence, and a single operational model.

Alexey Romanov

Head of Engineering

March 2026

Executive Summary

European mid-market companies are caught between two hard choices for their cloud infrastructure. They can stay on hyperscaler platforms and accept growing cloud resilience and sovereignty gaps as regulation tightens, or migrate to EU-native providers and sacrifice the ecosystem depth their engineering teams depend on. Neither approach can survive the reality of running cutting-edge, innovative production workloads under DORA, NIS2, and the EU Cloud Sovereignty Framework.

This white paper presents architecture patterns for a third approach: policy-driven workload placement across multiple providers, but governed through a single operational model. Rather than choosing between hyperscaler capability and sovereign compliance, organisations can define placement policies that place the right workload in the right environment, with continuous evidence that deployments always match placement policy.

~€2M

Target annual cloud spend for mid-market operations running production workloads across two to three providers, spanning Kubernetes clusters, data platforms, and AI/ML platforms

The patterns described here are designed for mid-market operations with production workloads spanning Kubernetes clusters, data platforms, and increasingly, AI/ML platforms. The approach inverts the traditional infrastructure decision: instead of choosing a provider and adapting workloads to its constraints, you define the governance requirements first and then place workloads in the environment that satisfies those requirements.

Disclaimer: This white paper is published by emma Technologies S.à r.l. and reflects the European regulatory landscape as of Q1 2026, including DORA, NIS2, and the EU Cloud Sovereignty Framework. It is intended as technical guidance for cloud architecture decisions and does not constitute legal, regulatory, or compliance advice. Organisations should consult qualified legal counsel for matters relating to their specific regulatory obligations.

1. The Architectural Problem

THE REGULATORY FORCING FUNCTION

The cloud infrastructure decisions that European mid-market companies made three to five years ago were primarily technical and economic. Which provider offered the best managed Kubernetes service? Where could the team get the most compute per euro? These decisions are now subject to regulatory scrutiny and compliance mandates that most of them were never designed to withstand. Three regulatory frameworks are converging to create this pressure:

DORA In application since January 2025, DORA requires financial entities and their ICT service providers to maintain comprehensive risk management frameworks, documented exit strategies for critical providers, and explicit concentration risk assessments. Article 28 mandates that organisations assess whether arrangements may reinforce ICT concentration risk. Article 29 requires preliminary assessment of that concentration, including whether a provider would be easily substitutable.

NIS2 With national transpositions now in force across the majority of EU member states, NIS2 extends cybersecurity risk management requirements to a far broader set of sectors — including cloud computing service providers, data centre operators, and any entity providing services to critical infrastructure. Fines for non-compliance can reach €10 million or 2% of global annual turnover. Senior management bears personal accountability for compliance.

EU CLOUD The EU Cloud Sovereignty Framework, operationalised through the European Commission's €180 million procurement tender in October 2025, introduces measurable sovereignty objectives across eight dimensions: strategic autonomy, legal jurisdiction, operational control, environmental sustainability, supply chain transparency, technological openness, security, and compliance with EU law. While currently applied to public procurement, the framework is envisioned as a reference point for the broader cloud market.

> *Your cloud architecture is now a compliance surface. Where workloads run, who controls the infrastructure, how you can prove policy enforcement, and whether you can exit a provider relationship — these are auditable questions with regulatory consequences.*

Why Binary Cloud Choices Fail

The instinct, when confronted with sovereignty requirements, is to simplify. Move everything to a single EU-native provider and eliminate the compliance ambiguity. Or stay on the hyperscaler, adopt their sovereign zone offering, and trust that their compliance investments will cover your obligations. Both approaches can fail under scrutiny.

The "EU-Only" Approach

Resolves jurisdictional questions but creates new problems. EU-native providers typically offer narrower service catalogues, less mature managed services, and smaller ecosystems. Migrating entirely often means rebuilding application architecture, retraining teams, and accepting operational trade-offs. It also creates its own concentration risk — a single EU provider is still a single provider.

The "Hyperscaler Sovereign Zone" Approach

Preserves ecosystem access but introduces governance complexity that is rarely transparent. Questions about control plane residency, personnel access, encryption key custody, and the applicability of non-EU legal instruments (the US CLOUD Act, in particular) are not always answered with the specificity that DORA and the EU Cloud Sovereignty Framework demand.

The Case for Policy-Driven Multi-Cloud

Sovereignty is not a property of a single provider. It is a property of the governance model that sits across providers. A policy-driven multi-cloud architecture inverts the traditional infrastructure decision with three structural advantages:

It eliminates the false binary

Workloads requiring strict EU sovereignty run on EU-native infrastructure. Workloads benefiting from hyperscaler ecosystem depth run on hyperscaler infrastructure. The governance layer ensures each workload runs where it should, with evidence that it does.

It reduces concentration risk by design

Rather than concentrating all workloads on a single provider and then attempting to write exit plans after the fact, multi-cloud governance distributes workloads across providers from the start. Exit capability becomes an inherent architectural property, not a contractual aspiration.

It produces audit evidence as a byproduct

When policy enforcement is embedded in the deployment pipeline, compliance evidence is generated continuously. It does not need to be assembled retrospectively before an audit. This aligns directly with DORA's requirement for documented, tested, and periodically reviewed governance frameworks.

2. The Define → Deploy → Govern Lifecycle

The architecture patterns in this white paper are organised around a three-phase governance lifecycle. Each phase produces a specific output that feeds the next, creating a closed loop between policy intent and operational evidence.

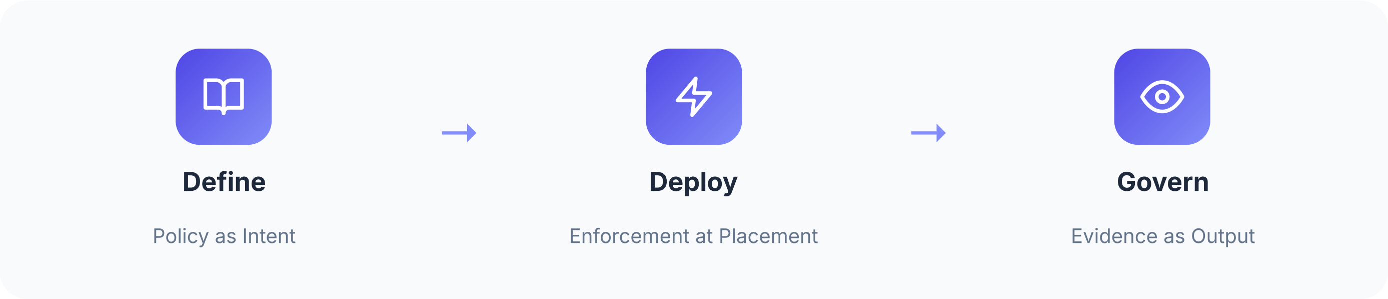

Define

Policy as Intent

→

Deploy

Enforcement at Placement

→

Govern

Evidence as Output

POLICY DIMENSIONS

Data Residency

Where data at rest and data in transit must be geographically located.

Operational Sovereignty

Who can access the control plane, management APIs, and support functions — and under which jurisdiction.

Encryption Key Custody

Whether encryption keys are managed by the provider, by the organisation, or by a third-party KMS within EU jurisdiction.

Exit Capability

Whether a tested migration path to an alternative provider exists for the workload.

Audit Evidence

What governance evidence must be continuously generated for the workload.

Workload Classification

Mapping each workload to a governance tier, and each tier to a set of policy constraints.

Workload Classification & Policy Enforcement

Not every workload requires the same level of governance. A development environment running non-sensitive test data has different sovereignty requirements than a production data platform processing customer financial records. The define phase includes a workload classification scheme that maps each workload to a governance tier, and each tier to a set of policy constraints.

TIER	REQUIREMENTS	INFRASTRUCTURE
SOVEREIGN-CRITICAL	Workloads processing personal data, regulated financial data, or data subject to specific residency requirements. Requires EU-resident operational access, customer-managed encryption keys, and tested exit capability.	EU-sovereign infrastructure only. Only IONOS, OVHcloud, Gcore or equivalent EU providers permitted.
GOVERNED-STANDARD	Production workloads without specific residency requirements but subject to general DORA/NIS2 compliance. Requires documented placement policy, provider risk assessment, and continuous governance evidence.	Hyperscaler or EU provider with documented risk assessment and governance evidence.
FLEXIBLE	Non-production workloads (development, testing, experimentation) with minimal governance constraints. Subject to cost and performance optimisation.	Any provider. Optimised for cost and developer experience.

DEPLOY: ENFORCEMENT AT PLACEMENT

In a policy-driven architecture, placement is evaluated against the workload's governance tier at deployment time. If a deployment request specifies infrastructure that does not satisfy the workload's policy constraints, the deployment is blocked — not after the fact through an audit finding, but at the point of intent.

Namespace-level policy labels

Each Kubernetes namespace is labelled with its governance tier. A namespace labelled sovereignty: critical can only schedule pods on nodes that satisfy sovereign-critical policy constraints.

Admission control for policy enforcement

A policy admission controller evaluates every deployment request against the namespace's governance tier before the Kubernetes scheduler places the pod. Violations are rejected with clear explanations.

Cross-cluster federation with policy awareness

Federation controllers manage workload distribution across clusters. Policy constraints determine which clusters are eligible — not just resource availability or latency, but governance compliance.

Govern: Evidence as Output

The govern phase closes the loop between policy intent and operational reality. The output is continuous governance evidence: a provable, auditable record that every running workload complies with its defined policies — not at a point in time, but continuously.

Why point-in-time audits fail: Traditional compliance approaches rely on periodic audits: once a quarter or once a year, an assessor reviews documentation, interviews stakeholders, and produces a findings report. This approach has two structural weaknesses. First, it produces evidence about a moment in time, not about the continuous state of operations. A workload that was compliant when the auditor checked may have drifted since. Second, it relies on documentation that may not reflect operational reality. A policy document that says "all customer data is processed in the EU" is not evidence — it is an assertion. Evidence is the operational data showing that it actually is.

Continuous governance: The govern phase replaces periodic assertion with continuous verification. The governance platform continuously evaluates every running workload against its policy tier and produces evidence records showing:

Policy Compliance Status

Is this workload currently running on infrastructure that satisfies all applicable policy constraints? Continuously evaluated against the governance tier.

Drift Detection

Has anything changed since deployment that affects policy compliance? For example, a provider configuration change, a new data flow, or a policy update that changes the constraints.

Evidence Chain

For each workload, a complete chain from the policy definition (intent) through the deployment decision (enforcement) to the current runtime state (proof).

Audit-Ready Packages

Evidence packages that can be exported for regulatory review, with the complete chain from policy definition through deployment decision to current runtime state.

Mapping Governance to Regulatory Requirements

The evidence produced by continuous governance maps directly to specific regulatory requirements. Organisations need audit-ready evidence packages that can be exported for regulatory review, with the complete chain from policy definition through deployment decision to current runtime state.

DORA 5–11 ICT risk management framework: Continuous policy compliance status demonstrates that the organisation has an operational — not just documented — ICT risk management framework for cloud infrastructure.

DORA 28 Third-party risk management: Deployment records showing policy evaluation against concentration risk constraints demonstrate that the organisation assesses and manages third-party risk at the point of each infrastructure decision, not only at contract review time.

NIS2 Art.21 Cybersecurity risk management measures: Evidence of policy-gated deployment and continuous compliance monitoring demonstrates implemented, operational risk management measures — not just planned ones.

EU CLOUD Sovereignty Framework objectives: Continuous evidence of data residency, operational control, and legal jurisdiction compliance maps to the Framework's sovereignty objectives, with the governance tier classification mapping to the Framework's assurance levels.

Reference Architecture: Mid-Market Hybrid Sovereign Operations

This section presents a reference architecture for a mid-market company running production workloads across a hyperscaler and an EU-sovereign provider, governed through emma's policy-driven placement model.

3.1 Architecture Overview

The reference architecture assumes an organisation with the following profile:

PRIMARY HYPERSCALER

AWS (EU region, e.g. eu-central-1 Frankfurt) for general-purpose compute, managed Kubernetes (EKS), data analytics, and AI/ML workloads.

EU-SOVEREIGN PROVIDER

A European-owned infrastructure provider for workloads requiring strict sovereignty — personal data processing, regulated financial operations, and workloads where the organisation needs to demonstrate full EU jurisdictional control.

GOVERNANCE LAYER

emma, operating as the policy engine and deployment orchestrator across both environments.

Network Topology

The two environments are connected via encrypted, dedicated interconnects (not public internet). Traffic between environments is subject to policy evaluation: data classified as sovereign-critical cannot traverse to the hyperscaler environment, and access from hyperscaler-resident services to sovereign infrastructure is controlled at the network level and logged for audit purposes.

Identity Federation

A single identity provider (IdP) manages access across both environments, with policy-based access controls that restrict which identities can access which environments. Sovereign-critical namespaces require identities with EU-resident attributes — meaning the human accessing the environment must be operating under EU jurisdiction, and the access must be logged with jurisdictional metadata.

Secrets Management

Encryption keys for sovereign-critical workloads are managed through a key management service (KMS) operating within EU-sovereign infrastructure, under the organisation's control. Keys for governed-standard workloads may use the hyperscaler's KMS with customer-managed keys. The governance layer tracks key custody as a policy dimension.

CI/CD Pipeline Design

The deployment pipeline is the enforcement point for policy. Every deployment request passes through emma's policy evaluation before reaching the target environment. The pipeline design includes:

1. **Build Stage**
Container images are built in a neutral CI environment and signed with a provenance attestation.

2. **Policy Evaluation**
The deployment request (specifying the target namespace, environment, and workload classification) is evaluated against the governance tier's policy constraints.

3. **Deployment**
If policy evaluation passes, the deployment proceeds to the target environment. If it fails, the deployment is blocked, and the failure reason is logged.

4. **Post-Deployment Verification**
The governance layer verifies that the deployed workload matches the approved deployment specification — catching any drift between the deployment request and the actual runtime state.

> ***emma's Position in the Stack:*** *emma operates as a governance and orchestration layer, not as a data proxy or traffic intermediary. Workloads run directly on the target infrastructure — emma does not sit in the data path, does not proxy API calls, and does not access customer data. The compute, storage, and networking remain native to the target provider. A governance layer that itself becomes a point of concentration or data access would undermine the sovereignty model it is meant to enforce.*

3.2 Workload Distribution Example

To illustrate how policy-driven placement works in practice, consider a mid-market financial services technology company with three categories of production workloads:

CUSTOMER DATA PLATFORM (SOVEREIGN-CRITICAL)

Processes personal financial data for EU customers. Data residency, operational sovereignty, and encryption key custody all require EU-sovereign infrastructure. The platform runs on the EU-sovereign provider, with access restricted to EU-resident identities and encryption keys managed through the organisation's EU-resident KMS. Exit capability is tested quarterly, with documented migration procedures to an alternative EU provider.

ANALYTICS AND ML PIPELINE (GOVERNED-STANDARD)

Processes anonymised, aggregated data for business intelligence and model training. No identifiable personal data is present after the anonymisation stage. This workload runs on the hyperscaler (AWS eu-central-1) for access to managed analytics services and GPU compute for model training. Policy constraints require documented provider risk assessment and continuous governance evidence, but do not require EU-sovereign infrastructure.

DEVELOPMENT AND STAGING ENVIRONMENTS (FLEXIBLE)

Non-production environments running synthetic data. These workloads run on the hyperscaler for cost optimisation and developer tooling integration. Minimal governance constraints — primarily cost management and access control.

The governance layer ensures that these placement decisions are not conventions but enforced policies. A deployment request that attempts to place a sovereign-critical workload on hyperscaler infrastructure is blocked. If components of a sovereign-classified data pipeline are provisioned on infrastructure that does not meet the required compliance tier, the deployment is flagged and prevented.

3.3 Operational Considerations

Cost Management

Multi-cloud operations introduce cost complexity. The governance layer provides unified cost visibility across both environments, with cost allocation mapped to workload classification. This cost overview also includes inter-cloud data transfer and egress costs, as they materially affect the true cost of portability and redundancy. It allows organisations to understand the cost premium of sovereign operations (which is real, but quantifiable) and make informed decisions about workload classification.

Team Skills

Mid-market organisations typically do not have dedicated platform teams for each cloud provider. The governance layer abstracts provider-specific complexity where possible, presenting a unified deployment and monitoring interface. Teams interact with the policy-driven deployment pipeline, not with provider-specific APIs directly. Provider-specific expertise is needed for infrastructure configuration and incident response, but day-to-day operations are normalised.

Incident Response

When incidents occur, the governance layer provides the context needed for response: which workloads are affected, what their governance classification is, which policies apply, and what the compliance implications of the incident are. This is essential for DORA's incident reporting requirements, which mandate classification and reporting of ICT-related incidents based on their impact on critical functions.

4. AI Workload Residency

AI workloads present a distinct governance challenge that is increasingly relevant for mid-market operations. The EU AI Act, in application since August 2025 with phased enforcement through 2027, introduces risk-based requirements for AI systems that directly affect cloud infrastructure governance.

THE TRAINING-INFERENCE SPLIT

Training Workloads

Consume large datasets, require significant GPU compute, and run for extended periods. The data used for training may include personal data, proprietary business data, or data subject to specific residency requirements. Where training runs — and what data it accesses — is a governance decision.

Inference Workloads

Serve predictions or generate outputs in real time. They may process personal data as inputs and produce outputs that themselves may be subject to governance requirements (an automated decision, a risk score). Can be optimised for latency and cost on hyperscaler infrastructure.

> *Policy-driven placement powered by emma allows organisations to govern training and inference workloads independently. Training that requires access to sovereign-critical data can run on EU-sovereign GPU infrastructure. Inference that serves non-sensitive outputs from a trained model can run on hyperscaler infrastructure for latency and cost optimisation.*

MODEL GOVERNANCE AND PORTABILITY

A model trained on one provider's GPU infrastructure must be deployable on another provider's inference infrastructure — otherwise, the governance layer has created a new form of lock-in. Model portability requires standardised model formats (ONNX, for example), provider-agnostic serving infrastructure, and governance metadata that travels with the model: where it was trained, what data it was trained on, and what governance tier applies to its inputs and outputs.

> *emma's governance model extends to AI workloads by treating models as governed artefacts with lifecycle policies. A model trained on sovereign-critical data can carry that classification throughout its lifecycle, including when it is deployed for inference on different infrastructure.*

5. Implementation Path

Adopting policy-driven multi-cloud governance is not an overnight transformation. For mid-market operations, we suggest a practical implementation path that is phased, allowing organisations to build governance capability incrementally, starting with the workloads that face the most immediate regulatory exposure and expanding from there. emma supports organisations through every phase.

PHASE 1

Assess and Classify

Inventory current workloads and data flows across all environments. Classify workloads by governance tier. Identify the workloads where sovereignty requirements are unambiguous (personal data processing, regulated functions) and where they are currently unaddressed. Map existing provider relationships against DORA concentration risk requirements.

4–6 weeks

PHASE 2

Establish the Governance Layer

Deploy emma as the policy engine and connect it to existing cloud providers. Define initial policies for the highest-priority governance tier (sovereign-critical). Integrate policy evaluation into the existing CI/CD pipeline. Establish baseline governance evidence generation and validate against audit requirements.

6–10 weeks

PHASE 3

Extend to Sovereign Infrastructure

Connect an EU-sovereign provider to the governance layer. Migrate sovereign-critical workloads to EU-sovereign infrastructure through policy-driven deployment. Validate exit capability by testing migration procedures. Configure identity federation, network topology, and encryption key custody across environments.

8–12 weeks

PHASE 4

Operationalise Continuous Governance

Activate continuous governance monitoring across all environments. Begin generating audit-ready evidence packages for regulatory review. Expand policy coverage to governed-standard workloads. Establish drift detection, incident response integration, and unified cost visibility across providers.

Ongoing

Conclusion

The regulatory environment for European cloud operations is no longer a future concern — it is a current reality. DORA is in enforcement. NIS2 transpositions are live across the majority of EU member states. The EU Cloud Sovereignty Framework has moved from policy document to procurement standard.

For mid-market companies, the architectural response cannot be "choose one provider and hope for the best." The complexity of the regulatory and innovation landscape demands a governance model that works across providers, enforces policy at the point of deployment, and produces continuous evidence that policy intent matches operational reality.

The define → deploy → govern lifecycle makes sovereignty an engineering pattern, not a compliance overlay

Sovereignty is a property of how you operate, not where you host

Multi-cloud governance reduces concentration risk as an inherent architectural property

Continuous compliance evidence replaces periodic audits with provable, operational governance

emma provides the governance layer that makes this pattern operational for organisations that need to act now.

For a board-level framing of this architectural approach, see the companion Executive Guide: "The Cloud Decision Your Board Will Ask About."

About emma

The cloud operations platform for distributed sovereign infrastructure

emma is built for organisations running distributed workloads across hybrid and multi-cloud environments — where operational complexity, unpredictable costs, and strict regulations get in the way of building.

It provides a single, policy-driven operating layer that spans hyperscalers, regional European cloud providers, AI-optimised infrastructure, and on-premises environments. From one dashboard, engineering, platform, and finance teams can deploy workloads, enforce compliance policies, and view and govern their entire infrastructure footprint. Sovereignty and compliance are built in through proactive guardrails, not bolted on.

With a vendor-neutral architecture, emma ensures organisations can choose the right infrastructure for cost, performance, and regulatory requirements — without lock-in or operational trade-offs.

Data Residency & Jurisdictional Control

Data stays within defined geographic and legal boundaries, enforced at the infrastructure layer across every environment rather than relying on policy documents.

Policy-Driven Workload Placement

Provider eligibility and sovereign constraints are maintained as infrastructure guardrails at deployment time, ensuring workloads run only in approved environments across providers.

Network Boundary Control

Cross-cloud data moves over emma's private networking backbone instead of the public internet, reducing cross-border public routing exposure within and between jurisdictions.

Continuous Governance & Compliance Evidence

Workloads are continuously evaluated against policies aligned with GDPR, NIS2 and DORA. Audit-ready logs mean compliance evidence is generated automatically.

Provider Portability & Exit Readiness

Tested migration paths across providers reduce concentration risk and meet regulatory exit capability requirements before they become urgent.

Encryption & Key Sovereignty

Encryption and key management policies are enforced at deployment time. Organisations retain control over who holds the keys, regardless of where workloads run.

2021
Founded & HQ in Luxembourg

15+
Cloud providers supported

~90
Cloud engineers at your service

NO
US Cloud Act applicability

CERTIFICATIONS & FRAMEWORKS

emma operates under internationally recognized security and compliance frameworks, including ISO-certified security management and SOC 2 Type II audited controls, with data protection aligned to GDPR and resilience aligned with NIS2 and DORA.

ISO/IEC 27001:2022 CERTIFIED AICPA SOC NIS2 DORA GDPR

emma

Ready to Operationalise Cloud Compliance?

See how emma's unified cloud operations platform turns regulatory complexity into infrastructure-level controls.

Book a Demo

emma

Make Cloud Work for You

emma Technologies S.à r.l. · 19-21, route d'Arlon · 8009 Strassen · Luxembourg

emma.ms