

ML LABS INTELLIGENCE

Building Multi-Tenant Isolation for Regulated SaaS

A regulated SaaS platform needed per-tenant auth, storage, and feature configuration — not just database partitioning. ML LABS built the isolation architecture.

CASE STUDY

Author Omar Trejo

Date 2026-03-03

ML LABS

mlabs.com/intel/multi-tenant-regulated-saas-platform

A regulated multi-tenant SaaS platform needed to serve organizations with fundamentally different requirements — different authentication methods, different integration endpoints, different storage boundaries, and different feature configurations — from a single codebase. Standard multi-tenancy patterns (shared database with tenant ID columns) could not handle the scope of per-tenant variation. Every layer of the stack needed organization-scoped configuration that enforced strict isolation without requiring separate deployments.

ML LABS designed and built the per-tenant isolation architecture across auth, storage, integration, and feature control. The system now onboards new organizations through scripted automation with drift detection, serves tenants with completely different auth flows from the same application, and has maintained zero cross-tenant data leakage in production.

The Problem

The platform served organizations across regulated industries. Each organization arrived with its own identity provider, its own compliance requirements, and its own integration surface. One tenant authenticated through OAuth machine-to-machine credentials bound to an EHR system. Another used SSO through an enterprise identity provider. A third used password-based auth with API keys for integration endpoints. Each needed different features enabled, different storage isolation guarantees, and different credential management.

The conventional multi-tenancy question — shared database or database-per-tenant — was the least interesting dimension. The real challenge was behavioral isolation: ensuring that tenant A's auth flow, feature flags, and integration behavior could not leak into tenant B's context, even when both shared the same compute infrastructure and codebase.

- Per-organization authentication methods ranging from machine-to-machine OAuth to SSO to API key auth

- Per-organization feature flags controlling integration behavior, enabled AI models, and workflow configuration
- Per-organization storage with dedicated cloud storage overrides for tenants requiring data isolation
- Cross-tenant safety requirements where a single misconfiguration could produce a compliance violation rather than an error message

Per-Organization Configuration as the Foundation

The isolation architecture started with a per-organization configuration system that controlled behavior at every layer. This went beyond typical feature flagging for gradual rollouts. Each organization's configuration governed which auth method was active, which integrations were enabled, which AI models could run, and which workflow behaviors were permitted.

The critical design decision was that misconfigured or missing configuration fails closed — blocking the request rather than falling through to a default that might expose another tenant's behavior. Getting this wrong in a regulated context means real consequences: clinical data sent to a system that was not expecting it, or one tenant silently receiving another tenant's workflow behavior.

The hardest bugs in multi-tenant regulated systems are not crashes. They are configuration leaks – one tenant silently receiving another tenant's behavior because a configuration was evaluated in the wrong scope.

Per-Organization Authentication

Each organization's authentication flow was bound at onboarding and enforced at every request. The platform supports multiple auth methods — machine-to-machine OAuth, SSO federation, API key auth — with each organization using the method that matches its infrastructure.

The critical safeguard is cross-tenant session isolation. A session established for Organization A cannot be carried into Organization B's context, even if the same user has access to multiple tenants. ML LABS removed all fallback mechanisms — there is no default organization. Every request requires explicit organization resolution, and ambiguity produces a rejection rather than a guess.

Inbound Request



Resolve Organization



Authenticate via
Org-Specific Method



Enforce
Cross-Tenant Isolation



Apply Per-Org
Configuration

Storage Isolation and Cross-Tenant Safety

Data isolation operated at multiple levels. Every data access path is scoped to an organization, and some organizations required dedicated storage for data residency or audit obligations. ML LABS built per-organization storage isolation that supports both shared and dedicated configurations depending on each tenant's compliance requirements.

The more dangerous surface was batch operations that touch multiple organizations' data in sequence. ML LABS enforced hard safety checks on every operation — an organization mismatch produces a failure, not a log entry. The most dangerous multi-tenant bugs are silent writes to the wrong tenant's data during operations that otherwise appear successful.

// **Zero cross-tenant data leakage is not a testing outcome. It is an architectural property — enforced by organization-scoped data access and safety checks that fail hard on mismatch.**

Onboarding Automation and Drift Detection

ML LABS automated the entire tenant onboarding sequence. Provisioning a new organization — with its auth method, configuration defaults, and storage resources — went from a manual multi-day process to a repeatable, idempotent scripted operation. The resulting configuration is validated against a known-good baseline before the tenant goes live.

The harder problem was configuration drift after onboarding. Infrastructure configuration drifts when engineers make manual changes during incident response. ML LABS built automated drift detection that compares the current state of organization-scoped resources against the declared configuration on a scheduled basis. Any deviation triggers an alert and blocks further provisioning until the drift is resolved. In regulated environments, a drifted configuration is a compliance finding, not just a technical debt item — catching it proactively matters.

Boundary Condition

This architecture carries overhead that is justified when tenants are organizationally distinct entities with different compliance requirements. If tenants are lightweight subdivisions of a single organization — departments, teams, project groups — per-tenant auth polymorphism and storage isolation are usually unjustified. A simpler role-based access model within a shared tenant boundary is the better fit. The full pattern applies when tenants span different identity providers, different regulatory regimes, or different integration ecosystems.

The difference between a first attempt at regulated multi-tenancy and an experienced execution is not incremental — it is categorical. The failure modes are silent (configuration leaks, cross-tenant session carryover, batch writes to the wrong org's storage) rather than loud. Pattern recognition from prior implementations compresses what would otherwise be an expensive learning curve, because the cost of discovering these failure modes in production is a compliance incident, not a bug report.

First Steps

1. **Audit where tenants currently differ.** List every dimension of variation: auth method, integration path, storage location, feature behavior, compliance constraints. If more than three dimensions vary across tenants, per-tenant configuration is not optional — it is the baseline.
2. **Implement org-scoped configuration with fail-closed evaluation.** Ensure missing or misconfigured settings block the request rather than falling through to defaults that might expose another tenant's behavior.
3. **Script onboarding and schedule drift detection.** The investment in automation pays for itself at the third tenant. The investment in drift detection pays for itself at the first compliance audit.

Practical Solution Pattern

Build regulated multi-tenancy as a per-organization configuration system that controls authentication, storage, integrations, and feature behavior — not just data partitioning. Ensure each tenant uses its own identity flow while sharing the same platform. Remove fallback mechanisms entirely so ambiguity produces rejection rather than a guess. Automate onboarding end-to-end and run continuous drift detection against the declared baseline.

This works because the failure modes in regulated multi-tenancy are silent. Configuration leaks, cross-tenant session carryover, and batch operations that write to the wrong organization's storage do not produce error messages — they produce compliance violations that surface during audits or during incidents. The architecture described here makes those failures structurally impossible rather than relying on careful coding to prevent them. The measure of success is not the elegance of the isolation model — it is whether the system reaches production with

zero cross-tenant leakage and onboards new organizations through automation rather than manual checklists. If the multi-tenant architecture is already live and the next risk is configuration drift or auth boundary gaps, an **AI Technical Assessment** can map the exposure before it becomes an incident.

References

1. Joint Task Force. **Security and Privacy Controls for Information Systems and Organizations (SP 800-53 Rev. 5)**. *NIST*, 2020.
2. Amazon Web Services. **Tenant Isolation - SaaS Architecture Fundamentals**. *AWS Whitepapers*, 2024.
3. Downs, J., Scott-Raynsford, D., & Vladimirskiy, A. **Architectural Considerations for Identity in a Multitenant Solution**. *Azure Architecture Center*, 2025.
4. Amazon Web Services. **SaaS Tenant Isolation Strategies**. *AWS Whitepapers*, 2024.
5. Auth0. **How to Choose the Right Authorization Model for Your Multi-Tenant SaaS Application**. *Auth0 Blog*, 2024.
6. HL7 International. **FHIR R4 Specification**. *HL7 International*, 2024.

Isolation

