

REV Spain

- Technical Architecture
- Diagrams
- Functional Overview

Technical Architecture

See the [Diagrams page](#) for the deployment diagram. See the [Functional Overview](#) for goals, context, and users.

1. Building Block View

REV Spain is a monolith. Main internal modules:

- **Intake module** — claimant onboarding and document/data submission entry point
- **Extraction engine** — parses ITA and RNT documents (fixed-structure official Spanish government PDFs) using direct PDF text extraction (no OCR, since layout is standardized)
- **Unified.to adapter** — alternative ingestion path fetching employee/salary data directly from client HR/payroll systems
- **Project & employee allocation module** — gathers project information and maps employees to projects
- **Eligibility rules engine** — evaluates R&D tax credit eligibility criteria
- **Report generator** — produces both financial and technical claim reports
- **Access control** — role-based access for claimants, reviewers, and admins

2. Deployment View

Stack

- Frontend: Vue.js 3 SPA
- Backend: Node.js / Express API
- Data & document storage: MongoDB, including GridFS for document storage (no separate object storage service)

Topology

- **VM 1** — runs the full application stack via Docker Compose (frontend, API, MongoDB)
- **VM 2** — reverse proxy handling the domain name and TLS termination, forwarding to VM

This is a single-server deployment: no horizontal scaling or failover across app instances currently exists.

3. Risks & Technical Debt

- **Single point of failure** — the entire application stack (including the database) runs on one VM; no redundancy or failover.
- **No documented scaling path** — current architecture sits in tension with quality goals like fast data gathering and claims centralization at higher volume.
- **Shared storage engine** — MongoDB serves both transactional data and GridFS document storage; document I/O could contend with regular query load as usage grows.
- **External dependency** — the Unified.to ingestion path introduces availability dependency on a third party for one of the two data-gathering flows.
- **PII handling** — sensitive employee/salary data is processed directly by the core system in both ingestion paths; encryption, access control, and retention policies should be explicitly reviewed and documented.

Diagrams

Workflow — Claim Generation Flow

Upload ITA/RN Tor connect Unified.to Extraction /Normalization Project Info +Employee Allocation Eligibility
Evaluation Report Generation (Fin. + Tech.)

Technical Architecture — Deployment

Internet VM 2 — Reverse Proxy (domain name, TLS) VM 1 — Docker Compose Host Vue.js 3 SPA (frontend)
Node/Express API (backend) MongoDB+ GridFS (documents)

Functional Overview

See the [Diagrams page](#) for the visual workflow.

1. Introduction & Goals

REV Spain lets companies self-manage the process of claiming Spanish R&D tax credits end to end: intake, eligibility check, extraction of employee/salary data from official documents, project information gathering, employee-to-project allocation, and generation of both financial and technical claim reports.

Quality Goals

- Self-service — minimal need for manual handling by internal teams
- Ease of use for non-technical claimants
- Fast data gathering (short time-to-claim)
- Centralization of claims across the organization

Stakeholders

- Top Management
- Head of Business Unit
- R&D Consultants

2. System Scope & Context

User Types

- Self-service claimants
- Internal reviewers
- Admins

External Systems

- **Unified.to** — third-party integration used as an alternative ingestion path when a client does not want to upload sensitive employee/salary data directly (e.g. pulls HR/payroll data from the client's own systems).

Sensitive Data Boundary

REV Spain handles sensitive employee and salary data directly in both ingestion paths: either after extraction from uploaded ITA/RNT documents, or via the Unified.to integration. There is no path that avoids the core system processing PII — this should be treated as a compliance-relevant boundary (encryption at rest/in transit, access control, retention policy).

3. Claim Generation Flow (Functional)

1. Claimant submits ITA/RNT documents, or connects via Unified.to
2. Employee & salary data is normalized
3. Claimant/reviewer provides project information
4. Employees are mapped to projects
5. Eligibility criteria are evaluated
6. Financial and technical reports are generated
7. Internal reviewer validates the claim before export/submission