PROJECT BRIEF

Humanitarian Jobs Dashboard

Real-Time Sector Employment Intelligence

Reference	PRJ-2025-001
Status	PRODUCTION
Date started	February 2025
Last updated	January 2026
Author	Jesus Baena
Live URL	https://baena.ai/demos/reliefjobs-dashboard
Source code	https://github.com/Jesus-Baena/2025-dashboard-reliefweb-jobs
Stack	n8n, Flowise, Metabase, Nuxt 3, PostgreSQL, Supabase, Docker Swarm, ReliefWeb API
Tags	AI, DataAnalysis, BusinessIntelligence

1. Background & Problem Statement

The year 2025 marked a turning point for the humanitarian sector. A convergence of funding freezes from major donor governments triggered what has been widely described as a “Great Contraction” — a rapid reduction in programme funding that has had immediate and measurable effects on the humanitarian workforce.

One of the most visible consequences has been the acceleration of precarity in humanitarian employment. Stable, long-term positions are increasingly being replaced by short-term consultancy roles — a structural shift toward gig-based engagement that threatens institutional memory, programme continuity, and the quality of aid delivery.

Despite the scale of this shift, there has been no real-time monitoring tool capable of tracking humanitarian workforce dynamics as they unfold. Standard analyses rely on retrospective reporting and often miss critical patterns: the high attrition rate among mid-career professionals, the gap between localization rhetoric and actual hiring practice, and the growing technical requirements being embedded in job descriptions without corresponding capacity building.

This project was built to fill that gap — a sovereign, data-driven intelligence platform that monitors the humanitarian job market continuously and surfaces the trends that matter to practitioners, researchers, and organizations navigating the contraction.

2. Goal & Objectives

Project Goal

To build a sovereign, AI-powered intelligence platform that monitors the ReliefWeb Jobs API to analyze the impact of the 2025 funding crisis on humanitarian workforce stability, localization trends, and technical skill requirements.

Objectives

#	Objective	Description
1	**Automated market monitoring**	Deploy n8n workflows to ingest daily job data from the ReliefWeb API, specifically tracking the ratio of consultancy vs. full-time roles to monitor sector precarity in real time.
2	**AI-driven extraction**	Leverage Flowise LLM extraction to read job descriptions and determine localization status (local vs. international hiring) and hidden technical requirements (SQL, KoBoToolbox, Power BI, GIS).
3	**Trend stabilization**	Implement SQL-based rolling seven-day totals to eliminate analytical noise from calendar cutoffs, ensuring an honest representation of market movements and downturns.
4	**Self-hosted visualization**	Utilize a containerized Metabase instance to serve high-density, interactive dashboards directly to stakeholders via a Nuxt frontend, with no dependency on third-party BI platforms.

3. Target Users

User Group	Description
Humanitarian professionals	Practitioners navigating a shrinking job market who need data on which roles, skill sets, and geographies remain resilient — and where the sector is shifting toward consultancy-based engagement.
Sector researchers	Analysts and academics studying workforce trends, brain drain, localization effectiveness, and the structural impact of funding contractions on humanitarian capacity.
NGOs & hiring organizations	Organizations seeking market intelligence on competitive compensation, skill demand shifts, and the effectiveness of their own localization and recruitment strategies.

4. Deliverables

Deliverable	Description	Status

Live intelligence dashboard	Metabase-powered UI displaying real-time job trends, hiring heatmaps, contract-type distributions, and experience-level breakdowns. Served via Nuxt frontend.	**Complete**
Automated data pipeline	n8n workflow ingesting daily job data from the ReliefWeb API, with automated rate-limit handling and error recovery to ensure reliable continuous collection.	**Complete**
AI extraction layer	Flowise-based LLM pipeline reading job descriptions to extract localization status and hidden technical skill requirements not captured in structured API fields.	**Complete**
Localization analysis	Data insights generated by AI extraction confirming the gap between localization rhetoric and actual local hiring patterns during the funding contraction.	**Complete**
Open-source pipeline	Publicly available GitHub repository containing SQL queries, n8n automation flows, and deployment documentation for the full stack.	**Complete**
Technical documentation	Beginner-friendly guide on deploying the n8n + Flowise + Metabase + Nuxt stack using Docker, including API configuration notes.	**In progress**

5. Technical Architecture

The platform is built as a sovereign, self-hosted data pipeline running on dedicated European infrastructure. The architecture was designed to be reproducible without enterprise tooling or cloud lock-in, using only open-source and self-hosted components.

Stack Overview

Layer	Technology	Role
Data ingestion	**n8n (self-hosted)**	Automated workflows polling the ReliefWeb API daily with rate-limit handling, batch pagination, and error recovery.
AI extraction	**Flowise + LLM**	Reads raw job descriptions to extract structured data not available in the API: localization status, technical skill requirements, contract classification.
Data warehouse	**PostgreSQL**	Central store for all ingested and enriched job data. Rolling seven-day aggregate views computed in SQL to eliminate calendar-cutoff noise.
Visualization	**Metabase (self-hosted)**	Interactive BI dashboards: trend lines, hiring heatmaps, contract-type

		distributions, experience-level breakdowns, geographic density maps.
Frontend	**Nuxt 3**	Public-facing embed layer serving dashboards with responsive layout and direct stakeholder access.
Infrastructure	**Docker (self-hosted)**	All components containerized and deployed on dedicated EU-based infrastructure. Fully sovereign — no external BI platform dependencies.

Design Decisions

- **Rolling totals over calendar periods:** The platform uses rolling seven-day aggregates rather than calendar-week or calendar-month totals. This eliminates the analytical noise caused by weekend posting patterns and month-boundary cutoffs, giving an honest view of market trends.
- **LLM extraction for hidden signals:** Job descriptions contain critical information (localization intent, technical skill requirements) that the ReliefWeb API does not expose as structured fields. LLM extraction runs against each description to surface this data systematically.
- **Sovereign data pipeline:** All data remains on self-hosted infrastructure. No third-party BI platforms or cloud analytics services are used. The visualization layer runs alongside the data warehouse, ensuring full control over data residency and access.
- **API resilience:** The ingestion workflow is designed to adapt to evolving API policies, including rate-limit changes and authentication requirements, without requiring a rebuild of the underlying pipeline.

6. Sustainability & Maintenance

The dashboard is designed to run indefinitely with minimal operational overhead. Sustainability considerations:

- **Near-zero operating cost:** All components run on self-hosted infrastructure at marginal cost. The ReliefWeb API is free. LLM calls are pay-per-use for new job descriptions only. Metabase has no licensing cost.
- **Automated daily operation:** The ingestion workflow runs daily without intervention. API changes (rate limits, parameter requirements) are the primary maintenance risk and are handled by adjusting workflow configuration rather than rebuilding the pipeline.
- **Open-source availability:** The full pipeline (SQL queries, n8n flows, deployment documentation) is published on GitHub, allowing the humanitarian community to fork, adapt, or build upon the work.
- **Modular and replaceable:** The architecture cleanly separates ingestion, enrichment, storage, and visualization. Any layer can be replaced independently — a different data source, a different LLM provider, a different BI tool — without affecting the rest of the pipeline.

7. Links & References

Live dashboard	https://baena.ai/demos/reliefjobs-dashboard
Project article	https://baena.ai/articles/jobs-relief
Source code	https://github.com/Jesus-Baena/2025-dashboard-reliefweb-jobs
LinkedIn post	https://www.linkedin.com/posts/jbaenanet_humanitariantech-dataanalysis-reliefweb-activity-7401125392348221440-vGZt
Portfolio	https://baena.ai

Jesus Baena | jesus@baena.ai | baena.ai | Kyiv, Ukraine