# The Refugees Urgent Support Hotline - section I

HERMES
HACKING FOR HUMAN RIGHTS

REFUGEES IN LIBYA
SEEKING ASYLUM IS A HUMAN RIGHT

**The report "The Refugees Urgent Support Hotline" was produced by Claudio Agosti for the Hermes Centre, as part of a collaboration with Refugees in Libya funded by Open Society Foundations.**

**June 2026**

# Introduction

The Refugees Urgent Support Hotline (RUSH) is a joint project by Hermes Center and Refugees in Libya (RiL).

Hermes Center is an Italian non-profit organisation established in 2012, founding member of the Rete per i Diritti Umani Digitali (Italian digital rights network) and member of the European Digital Rights coalition (EDRi).

Hermes Center's mission is to safeguard digital rights by raising awareness of both the risks and opportunities of digital technologies, and by promoting tools designed to ensure user safety, privacy, and freedom of expression.

Refugees in Libya (RiL) is a grassroots organisation made of refugees, asylum seekers and migrants. Founded in Tripoli in late 2021, following protests sparked by institutional neglect outside UNHCR offices, the organisation was formally registered in Italy in 2024.

Since 2023, RiL has operated a free WhatsApp hotline to collect reports of abuse, violence, arbitrary detention, and cases of vulnerability, as well as to provide useful information and refer victims to competent services.

The volunteers who manage the hotline have been relying on organically developed digital tools, which eventually became difficult to maintain and secure. Consequently, RiL reached out to Hermes for support in redesigning their case management system.

Our work began in January 2025, starting with the definition of a shared framework and the mapping of existing processes alongside partners such as Aequa (focusing on translation systems) and Liminal (providing OSINT and international litigation expertise). The partnership has so far sustained the project uniquely through its own resources.

Our goal is to build a secure pipeline integrating a WhatsApp Business number, a structured database and the targeted use of language models. This system could effectively support operators in triage, translation, and classification, while tracking reports over time to generate statistics for advocacy and build a robust evidentiary archive for strategic litigation.

This report provides an overview of the work carried out in establishing the RUSH pipeline. Replicability is a core objective of this project. We have highlighted technical specifications for those interested in deploying similar systems in other humanitarian contexts.

**Introduction** 3

**Part I**
**Open Source Tools and Technical Support in Humanitarian Crises** 5

Context of the RiL Hotline 5
Identifying the Challenges 6
Technical and Usability Challenges 6
Dictionary and Acronyms 6
The Value in the Data: Hotline and Archive 7
Selection and Integration of Open Source Solutions 7
Logical Architecture and Data Flows (RUSH) 7
Reproducibility of the RUSH Stack 8

**Part II**
**Using Language Models for Unstructured Data Management** 9

Evaluating Local Language Models 9
From the call to the aggregated data 10
Examples of extracted data 12
Challenges in Categorisation 14
Ethical Constraints 15
Data Governance and Metadata 16
Deployment Strategies and Physical Data Location 16
Threat Model Questions for Volunteers (Pre-hosting Assessment) 17
System Prompt for Automated Data Extraction 17

**Part III**
**Conclusion** 20

# Part I
# Open Source Tools and Technical Support in Humanitarian Crises

## CONTEXT OF THE RIL HOTLINE

During our initial assessment, we observed a high volume of calls to the RiL hotline. Conversations were recorded and saved on a "logbook" (technically a private forum). The lack of robust archiving processes made it increasingly difficult to search, navigate, and generate statistics over time.

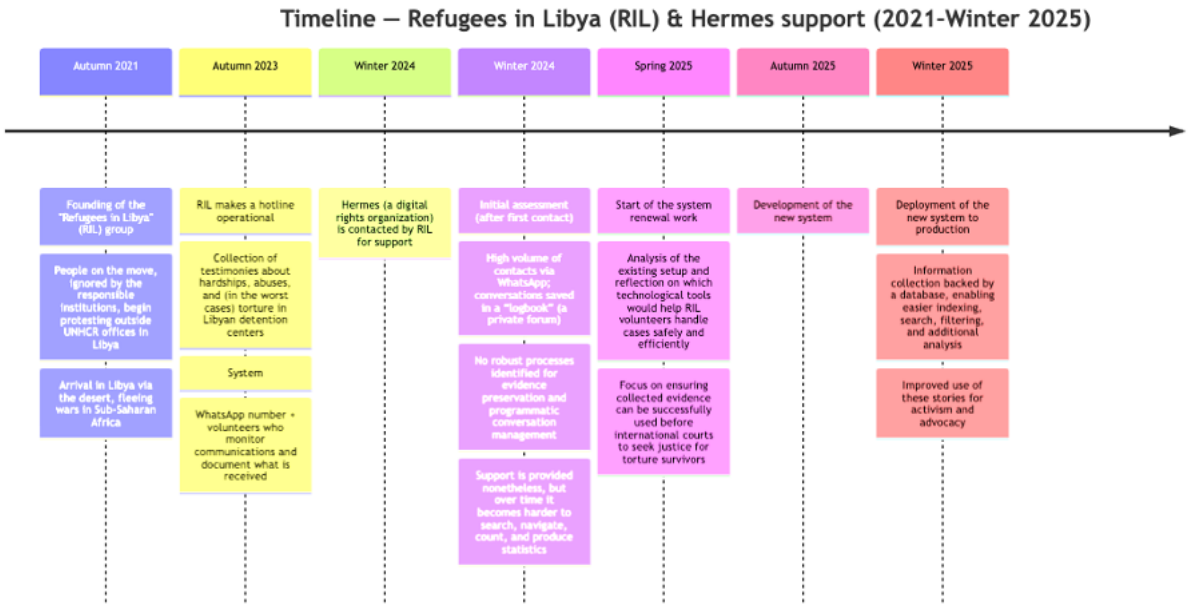


*Image 1: Timeline of RiL & Hermes support - showing the evolution from the Autumn 2021 protests to the deployment of the new system in Winter 2025.*

The infrastructure struggled to answer basic questions, such as: *How many calls were received? How many unique phone numbers have requested assistance? How many individuals have a UNHCR number?*

## IDENTIFYING THE CHALLENGES

RiL operates from the bottom up. The hotline emerged from the urgent need to save lives, prioritising immediate functionality over long-term efficiency and security.

### Technical and Usability Challenges

- **Inconsistent Interaction:** Unstructured interactions led to inconsistent data collection. Furthermore, using disparate translation tools (like ChatGPT) resulted in varying translations of Arabic names and places, complicating data matching.
- **Inadequate Archiving:** Using a private forum for case management failed to produce structured data, hindering complex statistical analysis.
- **Operational Bottlenecks:** The shared phone number was tied to mobile devices rather than a computer-automated system, limiting operations to one volunteer at a time. The lack of an integrated tagging system caused confusion regarding case urgency.
- **Internal Communication:** Internal notes regarding political updates or NGO closures were mixed with call transcripts, complicating operations.

## DICTIONARY AND ACRONYMS

- PIM (Person on the Move): Includes migrants, asylum seekers, and refugees moving across borders due to conflict, work, or safety.
- RiL (Refugees in Libya): A grassroots organisation defending the rights of people on the move in North Africa.
- RUSH (Rapid Urgent Support Hotline): The new technological pipeline designed to safely manage RiL's cases.
- LLM (Large Language Model): Artificial intelligence systems trained on vast amounts of text, used here to assist in translation and data structuring.
- OSINT (Open Source Intelligence): The collection and analysis of data gathered from open sources.
- GBV (Gender-Based Violence): Violence directed against a person because of their gender.

### THE VALUE IN THE DATA: HOTLINE AND ARCHIVE

A single unstructured message can capture the severity of an emergency, but an aggregated, structured dataset is essential to uncover migratory patterns and systemic abuses. An organised archive serves multiple high-impact goals:

- **Advocacy:** Grounding political communication in structured, verifiable data rather than isolated anecdotes.
- **Fundraising:** Demonstrating operational reach and justifying resource allocation to donors.
- **Continuous Support:** Tracking the evolution of specific cases over multiple interactions.
- **Strategic Litigation:** Transforming individual traumatic experiences into robust, legally viable evidence against specific authorities by contextualising and corroborating reports.
- **Accountability:** Evaluating the operational effectiveness of international bodies like UNHCR or NGOs.

### SELECTION AND INTEGRATION OF OPEN SOURCE SOLUTIONS

Together, we sought self-hosted, open-source applications that guarantee data sovereignty and allow integration with third-party services. We selected:

- **Chatwoot:** An open-source, multi-channel customer support platform to centralise hotline conversations – Supported Features on Channels - Chatwoot Developer Docs.
- **Baserow:** An open-source no-code database to build customised case repositories – Baserow Database API documentation.

### LOGICAL ARCHITECTURE AND DATA FLOWS (RUSH)

1. Channel: WhatsApp Business via Cloud API.

2. Operational Inbox: Chatwoot acts as the multi-channel front-desk. It separates channel access from physical phone ownership, simplifying shift management.

3. Operators use Tags, Internal Notes, and Assignments to manage the conversational timeline.

4. Structured Database: Baserow acts as the repository. Every relevant message/action is replicated here via webhooks/API to create a structured, queryable record.

5. Access Segmentation: Principle of least privilege is applied. Verification teams can access derived views with pseudonyms (data minimisation) without seeing full personal histories.

## REPRODUCIBILITY OF THE RUSH STACK

To support other organisations, RUSH is designed to be easily replicable.

### Minimum Hardware/Software Requirements:

- **Server:** 2–4 CPUs, 8 GB RAM recommended.
- **Disk:** SSD ≥ 100 GB (scaling with media attachments).
- **System:** Linux (e.g., Ubuntu LTS), Docker + Docker Compose.
- With a good design of questions, the needs of a **Local LLMs** might be reduced, but still might help to repeat some of the automated semantic extractions explained here. **It's Optional**, and should be maintained as a separate service on a dedicated machine; GPU recommended for lower latency.

# Part II
# Using Language Models for Unstructured Data Management

Conversation transcription yielded unstructured data. To generate structured data, we ran so-called “open source” local language models relying on a detailed "system prompt" to output a JSON object containing technical fields (phone number, name) and contextual fields (problem, location, authorities involved).

### EVALUATING LOCAL LANGUAGE MODELS

To ensure privacy and data sovereignty, we strictly avoided public cloud AI services, which pose risks regarding data retention, third-party access, and security filters that might suppress descriptions of trauma.

We tested four local, open-weight language models in parallel across 3,000 messages. The comparative approach mitigated the risk of a single model's biases “dominating” the dataset.

**Model Ranking & Performance:**

| Rank | Model | Producer | Features & Performance Notes |
|---|---|---|---|
| 1 | Qwen3 (14B, 6-bit quantised) | Alibaba Cloud | Best overall performance. Highly efficient due to quantisation, showing superior capability in handling diverse sub-Saharan Arabic dialects and extracting entities. |
| 2 | GPT-OSS (20B) | OpenAI | Open-weight Apache 2.0 model with a 128k context window. Provided very strong and consistent structured extraction. |
| 3 | DeepSeek | DeepSeek AI | Offers long-context processing up to 128k. Excellent reasoning, though slightly less consistent with the specific regional dialects compared to Qwen. |
| 4 | Llama 3 | Meta | Solid baseline model (8k context). Ranked lower in this specific experimental batch for Arabic dialect interpretation. |

## FROM THE CALL TO THE AGGREGATED DATA

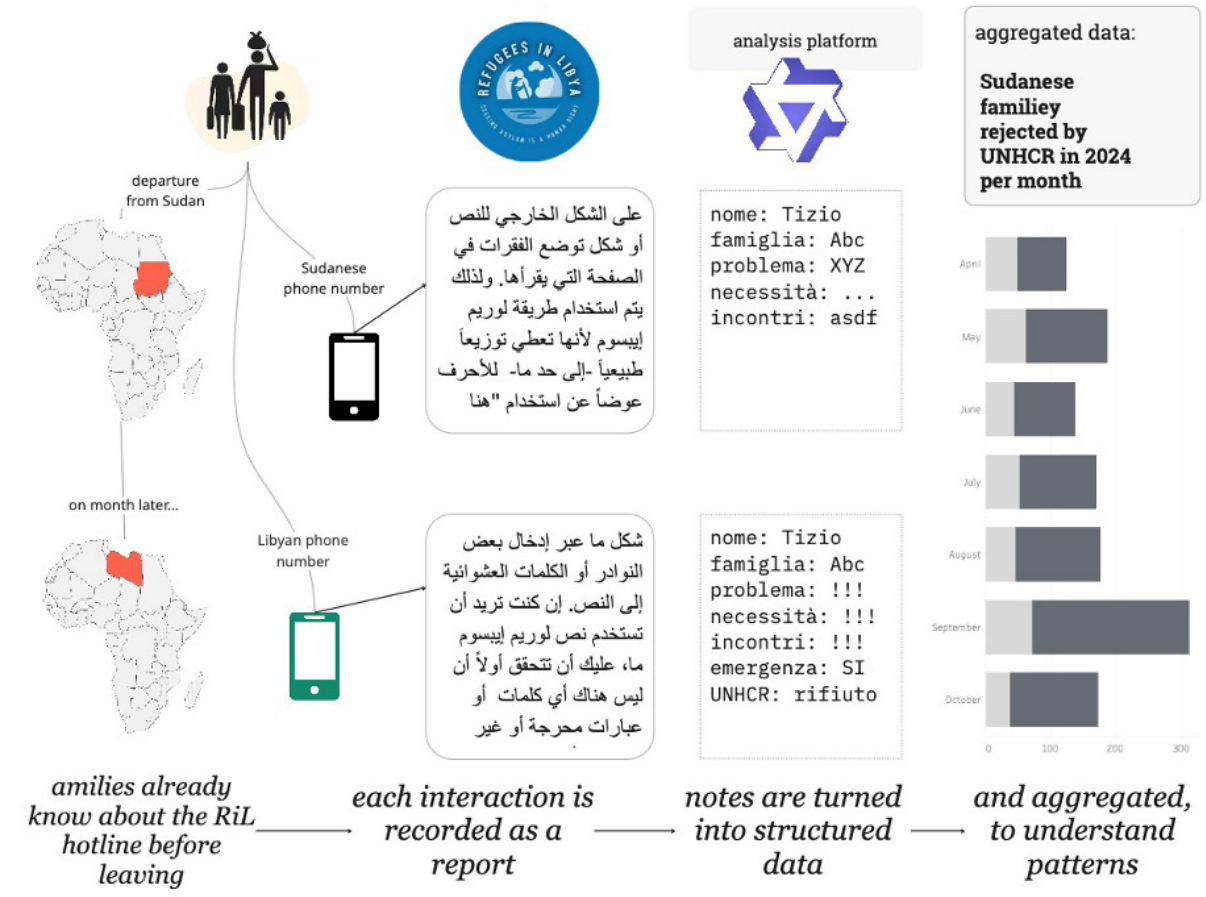


*The RUSH data pipeline mapping the journey of a family from Sudan to Libya. The workflow shows how initial awareness of the RiL hotline could lead to multiple unstructured text interactions over time, even if phone numbers change during transit. These individual reports are captured and processed by an analysis platform into structured data fields. Finally, the structured records are aggregated, in order to visualize mass patterns and identify critical trends, such as the monthly rate of Sudanese families rejected by the UNHCR in 2024.*

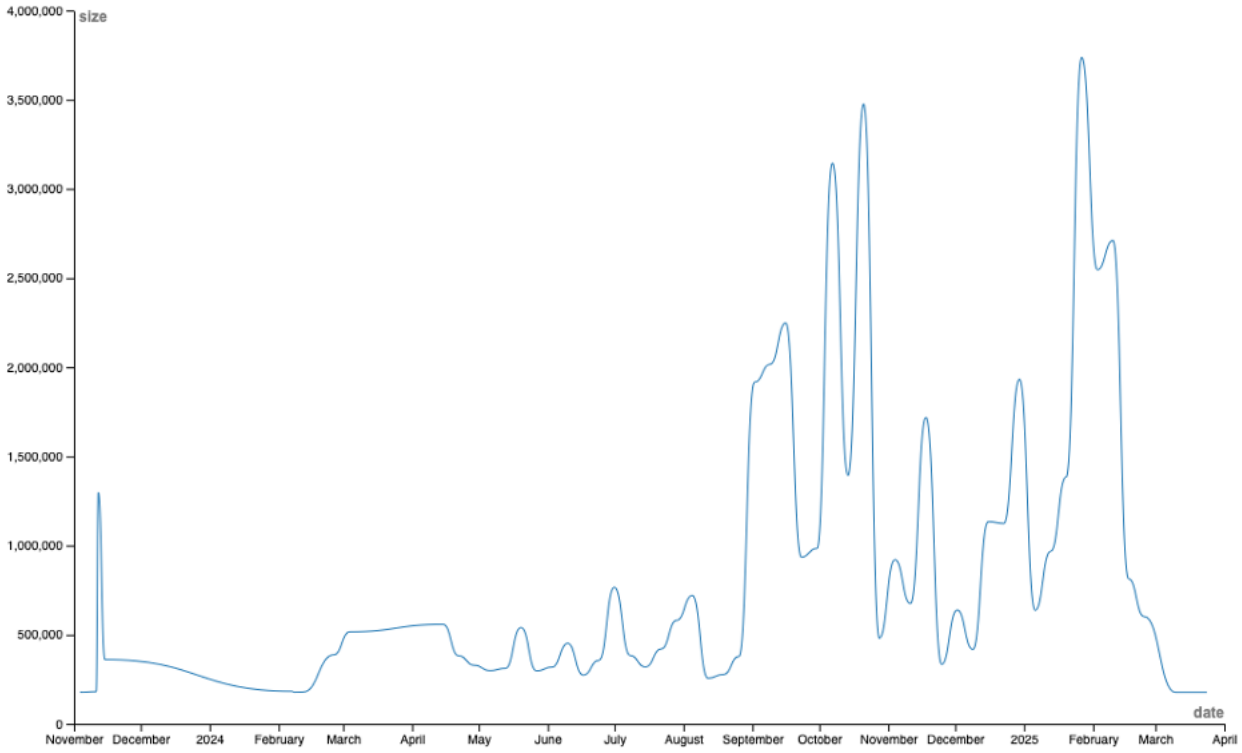


*Timeline showing the volume of messages received by the RiL hotline over one year and half. Peaks match specific regional events (such as opening new warzones) or increased visibility of RIL.*

We applied prompt chaining to transform narrative stories into machine-readable JSON formats.

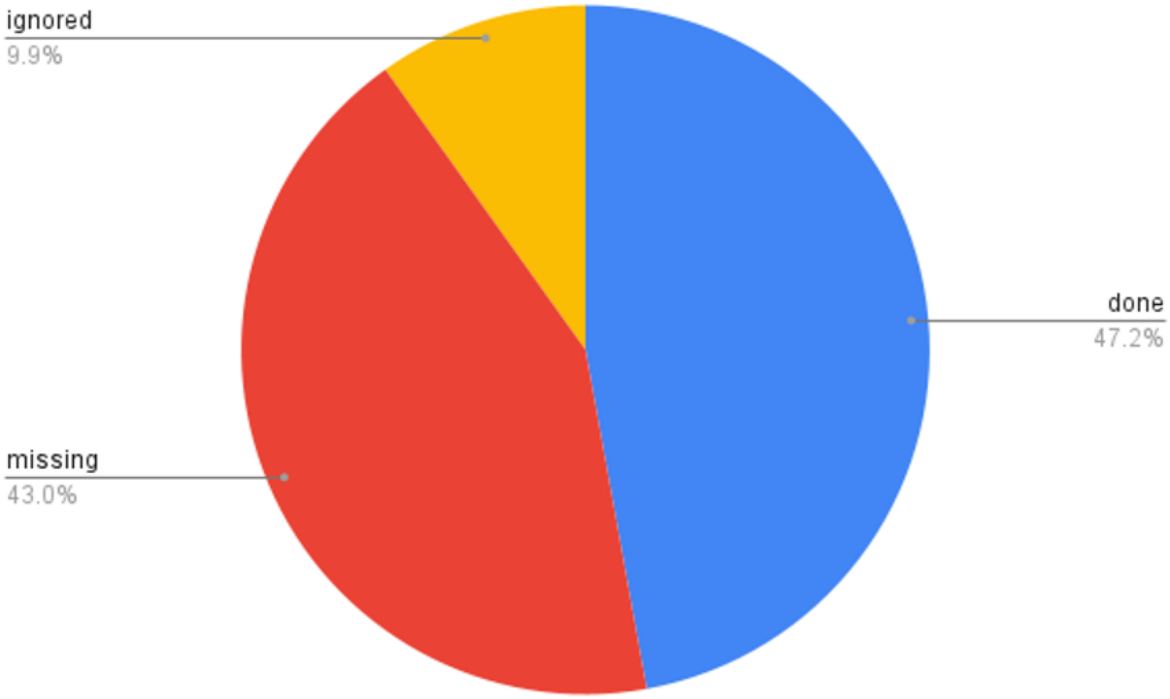


*Pie chart showing the breakdown of processed messages: 47% processed, 43% missing/excluded for this timeframe (November 2023 - April 2025), 9.9% ignored due to incompleteness.*

**Data Processing Insights & Limitations:** The process of parsing this data was instrumental in understanding the vast variance in cases, which directly informed how we designed tags and categories within Chatwoot.

### EXAMPLES OF EXTRACTED DATA


```
_id: ObjectId('6911a9c7a5e9e8af9bd4e9e0')
postId : "p243"
ExtraNotes : "The prisoner reports over 200 inmates, including Gambians, Guineans, N…"
Location : "Libya"
LocationDetails : "Asad prison, coordinates 32.0027126,10.8427731"
Name : "[redacted]"
OtherFamilyMembers : ""
PhoneNumber : "+218[redacted]"
Problem : "Detained in a non-governmental prison in Libya (Asad) with over 200 in…"
Request : "Voluntary repatriation to The Gambia and legal/humanitarian assistance…"
Summary : "A Gambian prisoner, identified as [redacted] is held in a non-gov…"
UNHCR : null
sudan : true
GBV : false
City : null
LocationNotes : "Asad prison coordinates 32.0027126 10.8427731"
Nation : "Libya"
Date : "Sat May 25 2024 09:34:54 GMT+0200 (Central European Summer Time)"
▾ AuthoritiesInvolved : Array (4)
    0: "Libyan Coast Guard"
    1: "Tunisian Coast Guard"
    2: "IOM"
    3: "European Border Authorities"
▾ AuthorityRoles : Object
    Libyan Coast Guard : "pushback"
    Tunisian Coast Guard : "pushback"
    IOM : "potential aid"
    European Border Authorities : "pushback"
```

```
_id: ObjectId('6911a9c7a5e9e8af9bd4ea01')
postId : "p188-7"
ExtraNotes : "File number mentioned (          ) but does not match UNHCR_CaseId…"
Location : "Sudan, Egypt, Libya"
LocationDetails : "Zamzam camp (Sudan), Tripoli (Libya)"
Name : "                    '
OtherFamilyMembers : ""
PhoneNumber : "+218          "
Problem : "Rape, severe beating causing eye injury, harassment, attempted rape, l…"
Request : "Protection and asylum-related reasons"
Summary : "Sudanese refugee from Darfur, Zamzam camp, subjected to rape, severe b…"
UNHCR : "                "
sudan : true
GBV : true
City : "Tripoli"
LocationNotes : " Zamzam camp () ()"
Nation : null
Date : "Mon Apr 15 2024 08:36:41 GMT+0200 (Central European Summer Time)"
▾ AuthoritiesInvolved : Array (3)
    0: "Libyan soldiers"
    1: "Tripoli commission"
    2: "Egyptian commission"
▾ AuthorityRoles : Array (3)
    0: "took phone, harassed"
    1: "harassed"
    2: "failed to provide protection or treatment"
```


*These images provide redacted examples of the structured data records generated from unstructured hotline conversations. Both samples display parsed JSON fields that categorise critical case details, such as problem descriptions, summaries, and specific geographical locations. Furthermore, the records demonstrate the use of boolean flags to identify specific case criteria like gender-based violence (GBV) , and show how the system extracts lists of involved authorities alongside their reported actions, ranging from harassment to pushbacks.*

However, it is crucial to note a significant analytical limitation: the models extracted data per interaction (each message or call session). In reality, a Person on the Move frequently contacts the hotline multiple times regarding the same issue.

To extract accurate statistics, these individual interactions must be aggregated and deduplicated using unique identifiers (like phone numbers or UNHCR IDs).

Releasing the raw extraction data at this stage would present a heavily biased and mathematically unbalanced view of the phenomena (e.g., counting the same event multiple times). Therefore, specific statistical findings have been omitted from this operational report.

**A thorough analysis of the aggregated data will be the subject of a forthcoming, dedicated report.**

### CHALLENGES IN CATEGORISATION

Even after refining the prompt, we encountered unpredictable model behaviour when dealing with authority roles.

For instance, mentions of the humanitarian organisation CESVI yielded highly inconsistent classifications. This table was generated by searching the keyword CESVI in a case-insensitive fashion.

| Authority | Role | Amount of times |
|---|---|---|
| CESVI | Helped | 1 |
| CESVI | contacted | 1 |
| CESVI | helped | 1 |
| CESVI | ignored 3 | 1 |
| CESVI | mentioned | 2 |
| CESVI | no response | 1 |
| CESVI | provided support | 1 |
| CESVI | providing hotlines | 1 |
| CESVI | received medical report | 1 |
| CESVI | responded with complaints only | 1 |
| CESVI | unknown | 1 |
| CESVI Organization | | 2 |
| CESVI Organization | BusyOnce | 1 |
| CESVI Organization | helped | 1 |
| CESVi | Protection request registered, no appointment set | 1 |
| Cesvi | Organization | 1 |
| Cesvi | denied entry | 1 |
| Cesvi | helped | 2 |
| Cesvi | helping | 1 |
| Cesvi | ignored | 1 |
| Cesvi | no response | 4 |
| Cesvi | receiving transferred persons | 1 |
| Cesvi | Detained | 1 |

*Table showing what is an inconsistent LLM output for the same institution, the Italian Cesvi Foundation. Beside the name, also the attributed behaviour ranges from "helped" and "contacted" to "ignored" and "responded with complaints only".*

This highlights an intrinsic limitation of relying on non-deterministic LLMs for synthesis: **without a strictly predefined vocabulary, the model generates fragmented and potentially biased interpretations.**

Moving forward, the prompt should be enriched with a pre-selected, restricted list of known authorities and permissible "roles" to ensure consistency.

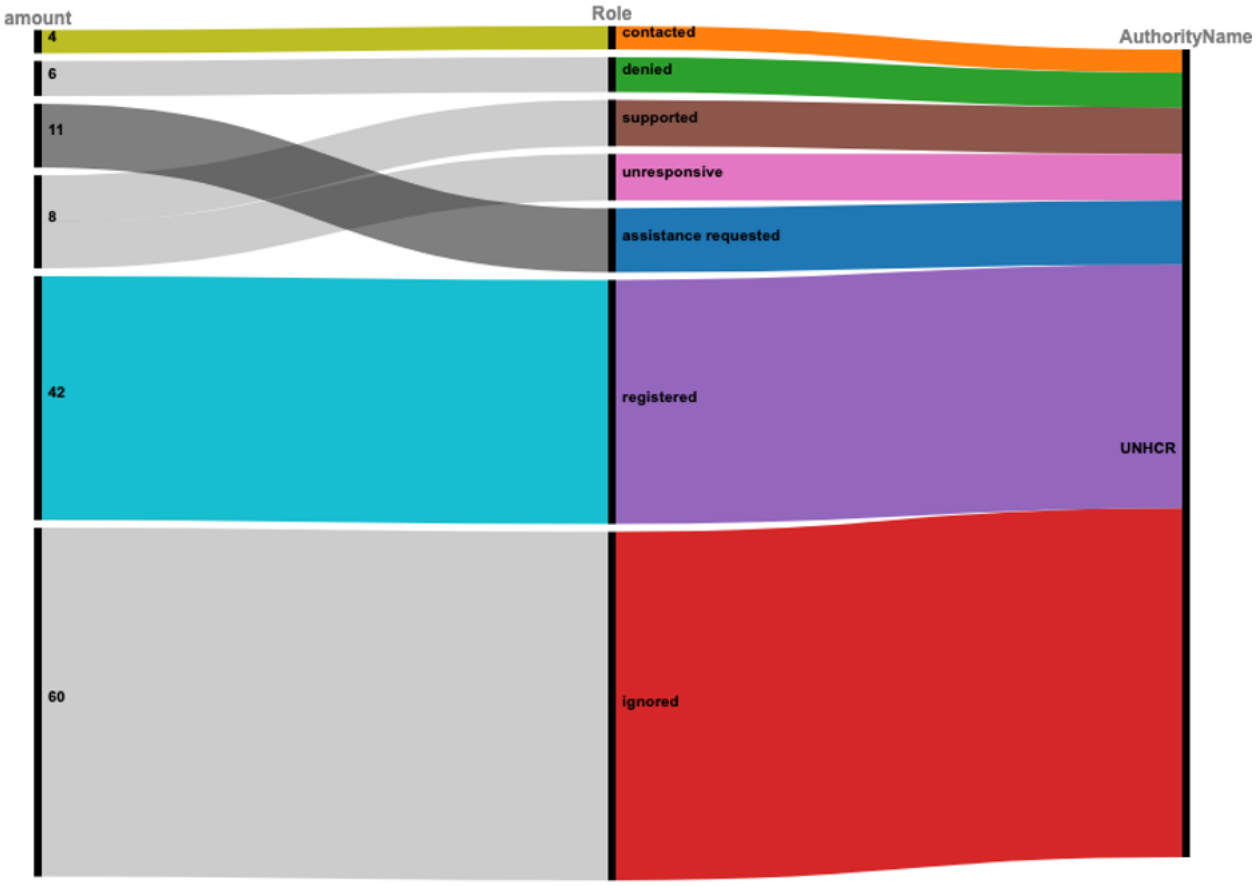


*Image: Sankey diagram attempting to map and aggregate the various roles and outcomes associated with UNHCR interactions*

## ETHICAL CONSTRAINTS

The use of AI in RUSH is treated as a matter of responsibility. A portion of these cases remains "uncomputable" – involving deep trauma and complex narratives that cannot be reduced to database fields.

The system operates on a dual track: some questions should stay free text (human-centric data collection), to allow people open expression and to be listened to fully. On the other hand, the two examples above display how much pattern analysis demands structured data, and this is solved by designing closed question, checkboxes and to train volunteers into learning a shared, standardized, vocabulary.

This also allows us to compromise with the "Do No Harm" principle. AI is strictly an assistive tool (human-in-the-loop), leaving all decision-making, empathy, and escalation in the hands of human operators.

### DATA GOVERNANCE AND METADATA

Every Person on the Move participates in one or more "conversations", but is associated with one, and only one, "contact" profile. The data provided can be textual, audio messages, videos, or images. It is crucial to save these contents and maintain them under the same case.

The new database would include:

- All received conversations in their original format.
- Their translations into a language accessible to the volunteers (English).
- Metadata and categorisation.

To add value to the collection and utilise it effectively across a multitude of cases, additional descriptive data (metadata) is required:

- **Technical Metadata (immutable):** The caller's phone number, full name, UNHCR number, region of departure, city and time of the call (noting there may be multiple calls per case), and languages spoken.
- **Assigned Metadata (selected by volunteers or assigned by the machine learning system):** Type of abuse, specific information regarding the abuse, and geographical location.
- **Process Metadata (internal states):** The internal sequence RiL uses to manage each individual case. Examples include: "received", "first contact", "in progress", "replied", "waiting", and "completed" (note: these statuses must be designed in collaboration with the volunteers implementing the process).

### DEPLOYMENT STRATEGIES AND PHYSICAL DATA LOCATION

When establishing a system like RUSH, the primary hosting options are:

- **Self-hosted:** Offers maximum control over data and configuration but requires strict discipline regarding backups, patching, and access management.
- **Public Cloud:** Easier to scale, but requires careful attention to administrative accounts, snapshots, logging, and contractual agreements.
- **Ethical/Solidarity Cloud:** Technically similar to public clouds, but preferable regarding governance and contractual clauses.

- **Hybrid:** Hosting the hotline (Chatwoot + Baserow) in the cloud while keeping the LLM local (or vice versa) to limit exposure and external dependencies.

**Criteria for deciding where the data "lives":**

- **Jurisdiction:** What access powers exist for authorities, and what real guarantees are in place?
- **Administration:** Who controls the servers, backups, and "root" credentials?
- **Copy Reduction:** How many replicas exist (local exports, attachments, caches) and for how long?
- **Exit/Migration Strategy:** The ability to export and restore data (DB dumps, CSV/API exports) without vendor lock-in.

### THREAT MODEL QUESTIONS FOR VOLUNTEERS (PRE-HOSTING ASSESSMENT)

Before finalising the hosting environment, the following questions must be addressed to outline the threat model:

1. What data must never be copied outside of RUSH (e.g., notes, exports, screenshots)?
2. From what devices would volunteers work (personal vs. shared)? Would encryption be enabled? Would Two-Factor Authentication (2FA) be active?
3. Who would be in charge of exporting data or downloading attachments? Who should be restricted from doing so?
4. What would the rapid response channel for incidents (e.g., a stolen phone, compromised account, or sending error) be?
5. What roles are genuinely necessary (e.g., triage, supervision, analysis), and what separation of permissions would be required?
6. What is the minimum acceptable retention period for logs, attachments, and the dataset?

### SYSTEM PROMPT FOR AUTOMATED DATA EXTRACTION

To standardise the unstructured narratives into machine-readable formats, we used a detailed system prompt. The story collected from the hotline has been appended to the end of this prompt to generate the extraction query.

```
You are a structured information extractor working on messages from a
hotline for refugees and asylum seekers.
Your response must be ONLY a valid JSON array.
Do not include any explanation, code, or text before or after.
Your goal is NOT to offer any help, you are here ONLY TO EXTRACT JSON.
If you CANNOT produce a valid JSON, return []. Nothing else.
Your task is to process each reported case, and extract detailed,
accurate, and consistent information into JSON format for downstream
case tracking.
All the values of the JSON format should be, where possible, translated
to English.

Each case comes with:
- A "Date": the date the report was created
- An "Author": the name of the volunteer who filed the message
- A "Message": free-text content describing the situation

---
**Fields to Extract:**

- UniqueId: A unique identifier for the person. This should be:
  - The PhoneNumber, if available;
  - If not available, the UNHCR_CaseId;
  - If neither are available, use null.
- PhoneNumber: A valid phone number. Acceptable formats:
  - Local or international: e.g., 091007704228, +218-91-7744022
  - May contain spaces, dashes, parentheses
  - Usually 8–13 digits, often starts with 0, +, or 00
- UNHCR_CaseId: A UNHCR case number, with these rules:
  - Starts with 5 and contains 11 or 12 digits
  - May appear as: 507-23-2784674, 507.25.4876605, 507232771167, UNHCR:
507232771167
  - May be preceded by text like UNHCR, case, ID
  - Do not confuse this with a phone number or a date
- Name: Name of the individual seeking help, if present
- OtherFamilyMembers: Names of other family members mentioned, if any
- Request: A practical request made by the person (e.g., relocation,
food support, legal aid)
- Problem: A brief description of the issue or situation (e.g., abuse,
detention, missing person)
- Date: Copy this directly from the Date field
- Volunteer: Name(s) of the volunteer(s). Use Author, plus any mentioned
in the message. Known names: Alessandro, Viola, David, Marc, Sophia.
- Location: City or region where the event occurred, if mentioned
- LocationDetails: Extra context (street, province, camp name, etc.)
- AuthoritiesInvolved: List of mentioned authority names (e.g., Police,
DCIM, Libyan Coast Guard)
- AuthorityRoles: 1-to-1 mapped list of each authority's role (e.g.,
detained, ignored, helped)
```

```
- Summary: A concise summary of the case
- Source: Source of the report (e.g., "Volunteer", "Direct report")
- Completeness: a number from 0 to 100 indicating how complete the case
is, based on the fields extracted
- ExtraNotes: Any other relevant information

---
**Critical Pattern Rules:**

- UNHCR_CaseId:
  - Must start with 5 and contain 11 or 12 digits
    - Valid formats include: 507-23-2784674, 507232771167, UNHCR
507.25.4876605
  - Look for keywords: UNHCR, case, ID
  - Never treat a phone number or date as a UNHCR ID
- PhoneNumber:
  - Usually 8–13 digits
  - Starts with 0, +, or 00
  - Formats include: 091007704228, +218-91-1234567, (091) 007704228
  - Never starts with 5 and has 10+ digits — that's likely a UNHCR ID

If a number is ambiguous:
- Extract it as a PhoneNumber only if it clearly matches phone format
AND is not preceded by UNHCR, ID, or case
- Never extract the same number as both PhoneNumber and UNHCR ID
- Text:
  - It should be only in English. If necessary, operate a translation.
```

# Part III
# Conclusion

As we conclude this initial phase, the new Chatwoot-based infrastructure is currently undergoing active testing by the RiL team.

We are optimistic that the extensive co-design process with the volunteers will be successful and result in meaningful improvements, ensuring that the new platform not only enhances technical efficiency and security but also genuinely supports the human complexity of the hotline's daily operations.

Partner organisations have demonstrated strong commitment to the project by contributing substantial internal expertise, staff time, and operational support to sustain ongoing activities. Additional external funding is now essential to implement the next phase of the RUSH project, focused on equipping hotline operators with the training, technical assistance, and operational support needed to strengthen and sustain frontline response capacities.

Regarding the data, the limitations we encountered during our first automated extraction have played a key role in shaping and refining our methodology.

Recognising the need to aggregate interactions by individual caller rather than by isolated message, **we have now launched a second, significantly more complex analysis.**

**We will publish a follow-up data report that will complement and complete the structural findings discussed here. Stay tuned for further updates at https://hermescenter.org.**