

GrieVoice/GemVoice: Multilingual Grievance- Intake Voice System

Technical note on design approach, language strategy, and pilot readiness for Gemini-based variant

[Download Full Technical Note](#)

[Schedule Demo](#)



Made with GAMMA

Why This Gemini-Based Version Exists

The Core Problem

Earlier Hume-based version demonstrated strong empathetic grievance intake in supported languages, but practical deployment contexts require broader multilingual coverage—especially in South African environments where English alone isn't sufficient.

This version optimizes for **language reach** rather than replacing the earlier concept. It accepts that conversational naturalness may require more deliberate refinement while gaining multilingual pathway access.

The Strategic Response

This Gemini-based variant addresses broader language inclusion as a core requirement, not a nice-to-have feature. It's a parallel path testing whether voice-and-messaging grievance intake can extend into South African languages beyond Hume's supported range.



How GemVoice Works: System Architecture

1

General Conversational AI Layer

Handles live exchange with callers through Gemini's multilingual capabilities

2

Custom Grievance-Intake Layer

Structures listening, questioning, information capture, and interaction closure

The grievance-intake layer defines conversation openings, reassurance protocols, when to let callers speak freely, how to handle sensitive issues, and what structured information to capture—including caller preferences, contact details, employer names, locations, dates, people involved, issue types, urgency levels, and desired outcomes.

Goal: Let callers speak naturally while capturing issues in ways downstream teams can actually use—without forcing rigid form-filling.

Structured Information Capture



Caller Details

Name or anonymous preference, contact information where provided



Organizational Context

Employer name, site or location, incident period or dates



People Involved

Individuals affected, others involved in the grievance



Issue Classification

Type, urgency level, desired outcome, free narrative description

Information is captured through natural conversation flow, not forced form-filling. The system asks clarifying questions when needed and handles sensitive topics with appropriate care.

Language Support Status

English: Strongest Performance

Most mature language with strongest end-to-end intake flow, question sequencing, and issue capture. Closest to pilot-ready baseline.

1

2

Afrikaans: Credible Viability

Not yet at English polish level but clearly usable enough to demonstrate meaningful intake capability and practical pilot potential.

3

Setswana: Functionally Testable

Can hold basic intake conversation, ask recognizable questions, capture key details, and complete flow. Not yet fully natural or production-grade.

4

Additional Languages: Pilot-Ready

Additional South African languages functionally testable at intake-flow level with refinement needed for naturalness and voice quality.

Key Distinction: "Working" means callers can understand questions, respond in relevant language, and enable accurate issue capture—not necessarily native-level polish yet.

Current Design Trade-Off

하 = 더미안-주
다-피오-거버거'S



Hume-Based Version

- More mature empathetic voice experience
- Stronger built-in conversational behavior
- Narrower language range
- Less refinement work required

Current Gemini Path

- Wider multilingual pathway
- Broader language reach
- More refinement work needed
- Production-grade requires tuning

The trade-off isn't between "possible" and "impossible"—it's between **narrower language support with stronger built-in conversational behavior** versus **broader language reach with more refinement work required** to reach the same smoothness.

Testing Results to Date

1

English Baseline

Strongest confirmed result: English works well as intake baseline

2

Afrikaans Flow

Credible intake conversation and complete labor grievance flow demonstrated

3

Setswana Test

Basic labor grievance intake flow functional with rougher naturalness

4

Multilingual Path

Not merely theoretical—demonstrably functional at pilot-oriented level

Testing focused on whether the system supports realistic grievance-intake interactions in different languages while maintaining coherent overall flow. Results show the multilingual path is functionally testable with refinement priorities identified.

Recommended Pilot Approach

01

Scope Selection

One grievance workflow (e.g., general labor grievances), English baseline, one or two additional South African languages

02

Deployment Configuration

Deployment for agreed workflow with active testing and refinement during pilot period

03

Support & Review

Weekly or fortnightly transcript/sample review to improve prompts, categorization, and language handling

04

Findings & Refinement

Pilot findings including recommendations for strengthening language handling and capture quality

Bounded pilot allows generating real evidence on caller comprehension, capture quality, categorization accuracy, language-specific drop-off points, and refinement priorities before broader rollout.



Cost Structure & What Drives Pricing

1

One-Time Setup Fee

Workflow design, deployment, testing, prompt configuration, and launch preparation

2

Monthly Platform Support

Hosting, monitoring, support, prompt refinement, bug handling, and pilot review work

3

Usage Allowance

Agreed pilot volume with overage handled separately if needed

Larger Cost Drivers

- Implementation effort
- Multilingual testing and refinement
- Workflow complexity
- Support expectations
- Reporting requirements

Not Just Infrastructure

Client pays for working, pilotable, multilingual grievance-intake system and refinement work needed to make it useful in practice—not just tokens or runtime.

Bottom Line: Next Steps

Present Strengths

- Strong English baseline performance
- Meaningful multilingual feasibility demonstrated
- Credible pilot framing established
- Clear accessibility rationale for broader language inclusion

Refinement Needed

- Additional language naturalness improvement
- Voice consistency across languages
- Structured implementation and governance

The Gemini-based variant is a **multilingual extension path** solving broader language inclusion that earlier stack couldn't address well. Right next step: bounded, well-scoped pilot testing real use, measuring capture quality, and improving language layers where they matter most.

Strategic Questions

- Where does multilingual intake add most value first?
- Which languages matter most for initial deployment?
- What quality level is sufficient for pilot legitimacy?