



MEERKAT LLC · VOICE AUTOMATION SUITE

small-business automation, built in Anna, Texas

Voice Automation Suite — start here

small-business automation, built in Anna, Texas

A voice intelligence layer for your Notion workspace, shipped as four importable n8n workflows. Speak into a webhook or call a phone number; transcripts, sentiment, prosody, and follow-up tasks land in Notion automatically. Urgent calls trigger Slack alerts.

What's in this folder

voice-automation-suite-v1.0.0/	
■■■ 00_START_HERE.md	← you are here
■■■ 01_WHAT_THIS_DOES.md	← one-page operator overview
■■■ index.html	← interactive package landing page
■■■ inflection-explorer.html	← real pitch contours for the 8 fixtures
■■■ docs/	
■ ■■■ HOW_IT_WORKS.md	← architecture + data flow narrative
■ ■■■ SOFTWARE_REQUIREMENTS.md	← every dependency, version, license, cost
■ ■■■ SETUP.md	← step-by-step installation
■ ■■■ INFLECTION_RULES.md	← how prosody → labels works
■ ■■■ CUSTOMIZATION.md	← swap providers, retune thresholds
■ ■■■ TROUBLESHOOTING.md	← common errors and fixes
■ ■■■ FAQ.md	
■■■ workflows/	← 4 importable n8n JSON files
■■■ notion/	← DB schema + 80 rows of seed data
■■■ voice-fixtures/	← 8 mood samples + extracted prosody.json
■■■ automation/	← .env template + setup orchestrator
■■■ assets/	← Meerkat marks
■■■ LICENSE.md	← end user license agreement
■■■ CHANGELOG.md	
■■■ VERSION	

The 5-minute path

1. Read [01 WHAT THIS DOES.md](#) — understand what you bought.
2. Open **docs/SOFTWARE_REQUIREMENTS.md** — confirm you have (or are willing to sign up for) each service.

3. Open **docs/SETUP.md** — fill in `automation/.env`, run `python3 automation/setup.py`.
4. Open `index.html` in any browser — visual map of the architecture and click-through to each workflow.
5. Open `inflection-explorer.html` in any browser — see real pitch contours and the urgency score for each of the 8 voice fixtures.

What we built

> Every voice memo, customer call, and dictated task lands in Notion as a structured row — with a transcript, a sentiment score, a prosody fingerprint, and an inflection-based urgency rating. Urgent callers get auto-escalated to Slack and a P0 task within seconds. Outbound TTS replies use a voice and tone matched to the priority of the message.

Versions covered

- n8n ≥ 1.50
- Notion API 2022-06-28
- OpenAI Whisper-1, GPT-4o-mini
- Deepgram Nova-3 (with prosody + sentiment)
- ElevenLabs multilingual_v2
- Twilio Programmable Voice
- AWS S3 (or any S3-compatible store)

Product info

- **Price:** \$25 USD — single-purchase digital download
- **Vendor:** Meerkat LLC — `meerkat-llc.com`
- **License:** see `LICENSE.md` — perpetual, internal-use, single-organization
- **Self-serve:** all four workflows install via the included `automation/setup.py`. No implementation hours included; documentation names the failure modes you'll actually hit.

Support

Self-serve. Read **docs/TROUBLESHOOTING.md** and **docs/FAQ.md** first; they cover the common errors. Direct questions to `support@meerkat-llc.com`. If you want a Done-for-You install or a multi-organization license, email `hello@meerkat-llc.com`.

"Meerkat LLC" and the Professor Meerkat mascot are marks of Meerkat LLC.

What This Does

A four-workflow n8n pipeline that turns voice into structured Notion data, with sentiment and prosody scoring on every message and automatic escalation when callers sound urgent.

The four workflows

1. Voice Capture → Transcribe → Inflection

Any audio (mobile recording, browser webhook upload, Slack huddle export, Granola memo) lands in your `Voice Captures` Notion DB with a Whisper transcript. In parallel, Deepgram extracts pitch, energy, speaking rate, pauses, sentiment, and emotion — all written to a linked `Inflection Analysis` row.

Use case: standup memos, sales call summaries, executive voicemails. Every voice note becomes searchable, sortable, and filterable.

2. Voice Command → Notion CRUD

Speak a command, the workflow transcribes it, classifies intent (create / update / query) with an LLM, and acts on your `Tasks` database. "Add a P0 task to send the SOC 2 packet to Acme by Friday" becomes a P0 task with the right title, status, priority, and source tag.

Use case: hands-free task capture for field teams, drive-time logging, voice-driven inbox triage.

3. Notion → TTS Voice Reply

Every five minutes, the workflow scans Notion for `P0 + Doing` tasks, generates a 2-sentence readback, picks a voice + style based on priority (urgent / professional / calm / empathetic), runs ElevenLabs with SSML inflection cues, stores the audio in S3, and logs a `Voice Replies` row.

Use case: morning task readbacks, daily plan summaries, "voice memos" for asynchronous teams, accessibility.

4. Twilio Call → Live Inflection → Urgency Alert

When a customer calls your Twilio number, the workflow records the call, runs full transcription + prosody + sentiment, scores urgency 0–100 against a rule matrix, and — if the score breaches threshold — auto-creates a P0 task and pings a Slack channel with the transcript and the recording URL.

Use case: customer escalation triage, sales hotline routing, support QA, churn early-warning.

What lives in Notion

Four databases, fully wired with relations:

- **Voice Captures** — every audio input, with transcript, audio URL, language, source, captured-at timestamp
- **Inflection Analysis** — 1:1 with Voice Captures; sentiment, emotion, pitch, energy, pause ratio, speaking rate, inflection pattern, confidence tone
- **Tasks** — voice-driven and manual tasks with status, priority, owner, due date, source, and a relation back to the originating Voice Capture
- **Voice Replies** — every TTS reply with the SSML used, the voice ID, the style, the audio URL, and links to both the originating capture and the related task

Each one ships seeded with 20 fake rows so the relations populate immediately and clients can demo without recording a single thing.

What's measured per voice capture

Feature	Range	What it tells you
Sentiment	-1.0 to +1.0	Tone polarity (positive/neutral/negative/mixed)
Emotion	label	joy, calm, anger, frustration, urgency, sadness, curiosity
Pitch mean	80–280 Hz	Voice register; high = excited or stressed
Pitch variance	5–80	Prosody movement; low = monotone, high = expressive
Speaking rate	80–220 wpm	Fast = urgent or anxious; slow = sad or careful
Energy (RMS)	0–1	Loudness; high = engaged or angry; low = tired or sad
Pause ratio	0–1	Silence over total; high = hesitation, stalling, sadness
Inflection pattern	label	rising, falling, monotone, varied
Confidence tone	label	confident, uncertain, tentative, assertive
Urgency score	0–100	Composite — drives Workflow 4's alert threshold

Why this is differentiated

Most "voice to text" pipelines stop at the transcript. This one preserves and acts on the **how it was said**, not just the words — which is where the actionable signal lives in customer support, sales, and exec workflows.

What you get for \$25

Every artifact in this folder. Workflows, schemas, seed data, scoring logic, 8 voice fixtures, automation script that does the entire deploy, and full documentation. No vendor lock — self-host, edit, extend, or migrate at will. One-time purchase, perpetual license (see [LICENSE.md](#)).

"Meerkat LLC" and the Professor Meerkat mascot are marks of Meerkat LLC.

© 2026 Meerkat LLC · [meerkat-llc.com](#) · [support@meerkat-llc.com](#) · *built with care in Anna, Texas*

Setup

End-to-end install. ~20 minutes if you have all credentials already; ~45 if you're signing up for services along the way.

Prerequisites checklist

Before you start, confirm:

- ☐ Self-hosted n8n ≥ 1.50 with a public HTTPS URL (or n8n Cloud)
- ☐ Notion workspace + admin access to create an integration
- ☐ OpenAI API account with billing enabled
- ☐ Deepgram account (free tier is fine for testing)
- ☐ (For Workflow 3) ElevenLabs account + S3 bucket
- ☐ (For Workflow 4) Twilio account + a phone number + Slack workspace
- ☐ Local: Python ≥ 3.9 , `pip install requests`

See `SOFTWARE_REQUIREMENTS.md` for full pricing and signup links.

Step 1 — Get your credentials

#	What	Where
1	n8n API key	Self-hosted n8n → Settings → n8n API → Create API Key
2	Notion integration token	https://www.notion.so/profile/integrations → New integration → "Internal" → copy the token
3	Notion parent page ID	In Notion, open the page where DBs should live → click ... → Connections → connect your integration. Copy the 32-char hex from the URL.
4	OpenAI API key	https://platform.openai.com/api-keys
5	Deepgram API key	https://console.deepgram.com → API Keys
6	ElevenLabs API key	https://elevenlabs.io → Profile → API Keys

#	What	Where
7	AWS access key + secret	IAM user with S3 PutObject on your bucket
8	Twilio Account SID + Auth Token	https://console.twilio.com (Dashboard)
9	Slack bot token	https://api.slack.com/apps → Create app → OAuth scope <code>chat:write</code> → install → copy <code>xoxb-...</code>

Step 2 — Configure `.env`

```
cd voice-automation-suite-v1.0.0/automation
cp .env.example .env
open -t .env # macOS – or use any editor
```

Paste each value into the matching variable. Save.

Step 3 — Run the deploy script

```
cd voice-automation-suite-v1.0.0
python3 -m pip install requests
python3 automation/setup.py
```

The script will:

1. Create the 4 Notion databases under your parent page (or reuse them if they already exist by name).
2. Wire the relations between them.
3. Seed 20 fake rows into each database (80 rows total).
4. Create credentials in n8n (Notion, OpenAI, Deepgram, ElevenLabs, AWS, Twilio Basic, Slack).
5. Patch each workflow JSON with the real DB IDs and credential IDs.
6. Import all 4 workflows into n8n.
7. Activate them.
8. Print the live webhook URLs and save Notion DB IDs to `.db-ids.json`.

It's re-runnable; if a database with the same name already exists under your parent page, the script reuses it.

Step 4 — (*Workflow 4 only*) Configure your Twilio number

In Twilio Console → Phone Numbers → your number → Voice Configuration:

- **A call comes in:** Webhook → `https://YOUR_N8N_HOST/webhook/twilio-voice` (HTTP POST)

Save. That's the only manual step.

Step 5 — Smoke test

```
./automation/test-workflows.sh
```

Fires the 8 mood voice fixtures at workflow 1 sequentially and shows compact JSON for each capture. You should see 8 new rows in **Voice Captures** and 8 linked rows in **Inflection Analysis** within ~30 seconds.

Then call your Twilio number. Speak normally, then hang up. Within 5 seconds you should see a new Voice Capture row with the recording URL and (if you sounded urgent) a Slack alert + a P0 task.

Common setup errors

"403 forbidden — page not shared" — Notion integration isn't connected to your parent page. Open the page → ... → Connections → add the integration.

"401 — Invalid API key" on n8n calls — double-check `N8N_API_KEY` and that your n8n version supports the public API (1.20+).

"404 — workflow type version not found" when importing — your n8n is older than the workflows expect. Upgrade to ≥ 1.50 .

Twilio webhook fails with 502 — n8n isn't publicly reachable. For local dev, use a tunnel (ngrok, Cloudflare Tunnel, or `n8n start --tunnel`). For production, expose n8n behind a real HTTPS domain.

See `TROUBLESHOOTING.md` for the long list.

Uninstalling

```
# Stop the workflows in n8n UI (or DELETE /api/v1/workflows/{id})
# Archive the 4 Notion databases (don't delete — preserves your historical rows)
# Revoke API keys at each provider's dashboard
```

"Meerkat LLC" and the Professor Meerkat mascot are marks of Meerkat LLC.

© 2026 Meerkat LLC · meerkat-llc.com · support@meerkat-llc.com · built with care in Anna, Texas

"Meerkat LLC" and the Professor Meerkat mascot are marks of Meerkat LLC.

