

Strata: multilingual media intelligence and rights data

From Webinson s.r.o., Košice, Slovakia

For Broadcasters, archives, media monitoring services and rights societies looking for multilingual speech, video understanding and rights-data technology

Date 2 October 2026 **Contact** ceo@webinson.com

1. Summary

Strata is a self-hosted service that puts what a broadcast programme contains on one timeline: what was said and by whom, in which language, what was on screen, what other systems reported about it, and the music rights data behind it. It is one program with an HTTP API of 74 operations, described in OpenAPI 3.1, and a console for people. Every model it uses is open and runs on the server: on Webinson's servers in Nuremberg, Germany, or in a partner's own environment.

Strata is developed by **Webinson s.r.o.**, a software company in Košice, Slovakia, with 15 years of software engineering behind it. Since spring 2023 Webinson has built and run **Vocabulift**, a language-learning product that works with speech in real video; Strata grows out of that work for media, monitoring and rights.

Our proposal: a pilot on your own radio and television material, in your languages, measured against your references, leading to a commercial agreement with technical assistance, a research and development cooperation agreement, or both.

2. What Strata does today

CAPABILITY	WHAT IT DOES	ENGINE
Speech recognition	43 languages, Greek among them. Every word with its start, end and confidence; sentences rebuilt with punctuation.	whisper.cpp with Whisper large-v3-turbo (MIT)

CAPABILITY	WHAT IT DOES	ENGINE
Speakers	Who spoke when, without enrolment: speaker labels on words, sentences and turns.	Silero VAD (MIT) and WeSpeaker ResNet34 (CC BY 4.0) on ONNX Runtime
Language of each sentence	Checked sentence by sentence, for programmes that switch, such as an English answer in a Greek bulletin.	Trigram statistics and script, with the language model's second opinion
Transcript correction	Your lexicons correct names and terms; each correction keeps the original.	Strata
Translation	Each sentence into up to five of 35 languages per analysis, keeping its timing.	EuroLLM-1.7B-Instruct (Apache 2.0) on llama.cpp; a larger model can take its place
Entities	People, organisations, places, events, works, products and dates; lexicon entries with their canonical names.	GLiNER multi v2.1 (Apache 2.0)
Topics	Stretches of the programme classified into your own taxonomy, without training.	mDeBERTa-v3 trained for multilingual inference (MIT)
Summary and keywords	A short summary and up to eight keywords, in the language you choose.	EuroLLM
Shots and scenes	Shot boundaries, hard cuts and fades, with shots of one place or conversation grouped into scenes.	Strata, on ffmpeg
On-screen text	Captions, tickers and credits with their position and confidence, for as long as they stay on screen. Greek and English today; other scripts by adding their data files.	Tesseract (Apache 2.0)
Events and correlation	Events from other systems, such as audio identification, playout, schedules or alerts, on the same timeline; for each event, the sentences, speakers, on-screen text, shots, entities and topics around it.	Strata
Rights catalogue	Recordings, works and parties with ISRC, ISWC, IPI name and base numbers and territories, validated with their check digits; bulk import and validation without storing.	Strata

CAPABILITY	WHAT IT DOES	ENGINE
Matching	Recording-to-work proposals with their evidence and confidence, across Greek and Latin script; people confirm or reject them, and every decision is kept.	Strata
Conflicts	Records that contradict each other, such as a link whose ISWC differs from the work's, found and resolved by people.	Strata
Cue sheets	Generated from identified music or imported from CSV, Excel or JSON; corrected with versions; submitted, approved or rejected; exported as CSV, JSON, Excel or XML, with their full history.	Strata

Example output

One event of a 44-second Greek news clip: an audio identification system reported a song, and Strata gathered what met it on the timeline. The ISRC leads to the catalogue's recording «Το φως της πόλης», written in Greek, and the bulletin's cue sheet is generated from this match. Trimmed:

```
{
  "event": {
    "kind": "content_match", "source": "fingerprinting",
    "start_ms": 33875, "end_ms": 43937,
    "data": { "confidence": 0.97, "recording": {
      "title": "Το φως της πόλης", "artists": ["Anna Demeti"],
      "isrc": "GRX9D2600001" } }
  },
  "speech": [
    { "text": "Καλή σας νύχτα.", "start_ms": 32170, "end_ms": 33330, "gap_ms": 545 }
  ],
  "on_screen_text": [
    { "text": "ΜΟΥΣΙΚΗ", "start_ms": 33875, "end_ms": 43937, "overlap_ms": 10062 },
    { "text": "Το φως της πόλης Άννα Δεμέτη", "start_ms": 33875, "end_ms": 43937, "overlap_ms": 10062 }
  ],
  "shots": [4]
}
```

3. API and integration

- **74 operations in ten areas**, described in OpenAPI 3.1, with an interactive reference at vocalift.com/docs and a functional guide with workflows and real examples.
- **Keys with scopes** for media, rights data, cue sheets and administration. Errors are RFC 9457 problem details whose pointers name the field; lists page with cursors.
- **Background work** accepts an Idempotency-Key, so retries never start it twice; job progress streams as server-sent events.

- **Webhooks** when media is ready, a job ends, a link is proposed, a conflict opens or a cue sheet moves on, signed as the Standard Webhooks specification describes (HMAC-SHA256) and retried until delivered.
- **Results** as JSON, layer by layer or as one merged timeline; transcripts and subtitles as VTT, SRT and text; CSV and JSON exports; cue sheets as CSV, JSON, Excel and XML.

4. Performance

We do not quote general accuracy figures; they say little about your programmes. In the pilot we report, on your material and against your references:

- word and character error rate by language, channel and programme type;
- word-timing deviation from reference timestamps, and speaker attribution;
- on-screen text accuracy, and the precision and recall of recording-to-work matching;
- processing speed (real-time factor) and cost per hour for each deployment option;
- the share of results below an agreed confidence threshold, for human review.

For orientation: the 44-second sample goes through all eleven stages in about 20 seconds on a laptop with Apple silicon, and in about two minutes on the shared, CPU-only server of the public instance. A GPU makes speech recognition many times faster than real time; hardware for your volumes is sized in the pilot.

5. Computing and deployment

Strata is one program, written in Rust, with one SQLite database file and model files pinned by their SHA-256 checksums. It runs on Linux, on CPU or with a GPU.

- **Hosted in the EU.** The public instance runs on Webinson's servers in Nuremberg, Germany, declared with NixOS: models are fetched by their checksums, loaded while analyses need them and released when idle, the language model starts on demand, and the database is copied every night.
- **Private deployment.** The same program inside your environment or a European data centre you choose. No media is sent to an outside AI service, and every model can be replaced, the language model by a larger one.

6. Cooperation, licensing and IP

Commercial agreement with technical assistance

A licence to Strata, run by Webinson or in your environment, with Webinson's documentation, interfaces and support for integration, configuration, testing and validation. Licence terms, deployment requirements, performance parameters and any restrictions on use or data processing are set out before integration.

Research and development cooperation agreement

Joint development where your material and needs go beyond what Strata does today: tuning for your languages and channels, logos, objects and actions in the picture, the same speaker recognised across

programmes, and your partners' exchange formats such as CWR and DDEX, validated in defined pilot applications on your platform.

Each party keeps ownership of its pre-existing intellectual property. Access rights and ownership of jointly developed results are agreed for each cooperation.

7. Data protection

- Processing under a data processing agreement in line with GDPR.
- Processing in the European Union, or in your own environment.
- Your material is used for your project only, not to train general models, and is deleted when the pilot ends unless agreed otherwise.
- Access limited to the named engineers working on the pilot.

8. Proposed pilot

WEEKS	WHAT HAPPENS
1–2	A sample set of radio and television with reference transcripts and rights data; success criteria and the deployment option agreed; API keys issued.
3–5	Your material analysed; lexicons and taxonomies built with you; first measurements.
6–7	Integration with a test instance of your platform through the API and webhooks.
8	Validation report: measured performance, computing needs, cost and the next step.

The timeline is indicative and is fixed together with you.

9. Company

Webinson s.r.o., Alžbetina 16, 040 01 Košice, Slovakia · webinson.com · ceo@webinson.com

15 years of software engineering. Webinson has developed Vocabulift since January 2023 and brought it to market in spring 2023. Other products include PitchPlace (campsite websites in 28 countries read into structured, translated records and matched against other sources), PitchPace (market intelligence), Spotibo (SEO tools) and Colorko (a creative platform).

Strata technology profile, version 2.0, 2 October 2026. Prepared by Webinson s.r.o. Online at vocabulift.com/partnership.