



Spreetropol Systems

Emergence

A governance layer for interchangeable language models.

Same question, different answer

Today, LLM answers cannot be verified — no provenance, no repeatability, no trace.

- No provenance
- No trace
- No control

Question · Which database fits our project?

Go with PostgreSQL.
Relational, mature — and your data is clearly structured.

Question · Which database fits our project?


Go with MongoDB.
Schemaless, quick to start — your data model is still changing.

Both plausible. Neither says where it came from.

Not the prompt. The layer.

Emergence is not a chatbot with a better prompt, but an auditable governance layer over interchangeable models — every answer has a traceable provenance.



- Model interchangeable
- Behavior configured
- Provenance traceable



The model changes. The layer stays.

Four phases around every answer

Every request passes through the same four phases — understand, reason from perspectives, generate, check — and the actual run can be inspected as a graph.



The trace shows a real turn as a graph — from classification to the gate.

One rule, visibly in effect

Behavior is configured, not talked into the prompt: same request, one rule added, a visibly different answer.

- Trigger + action
- Effective immediately
- Configuration, not prompt

47.2k tok1.8%

Memory aktiv

We are Nordlicht GmbH, a small Hamburg-based company making outdoor and everyday products. Our new product is called PocketPress. It is a hand-powered, portable espresso maker priced at 89 ?. You press the espresso by hand ? no electricity, no capsules. It fits in a jacket pocket. Our audience is commuters, campers, and office workers who dislike the coffee at work. Create a Facebook advertisement post for PocketPress. Only give me the post, nothing else.

DONE20.2 s0.40 · eff 0.27

Gedankengang

PocketPress — real espresso. No plug. No capsules. No excuses.

The coffee at your office is not coffee. It's brown water with ambition.

PocketPress is a hand-powered espresso maker that fits in your jacket pocket. You add ground coffee, you add hot water, you press — and you get a proper shot with real crema. No electricity. No capsules. No queue at the café on the corner.

☕ Fits in a jacket pocket
👉 Hand-pressed — works anywhere
🌿 Zero capsules, zero waste
🏠 Built for the 7:14 train, the campsite, and the 3 p.m. slump at your desk

Made in Hamburg by people who got tired of bad coffee in good places.

PocketPress — 89 €. Your espresso goes where you go.

👉 Order now at nordlicht-pocketpress.de

WHEN (TRIGGER)

session.mode>=1

THEN (ACTION)

When the user requests advertising or paid social copy (a Facebook/Instagram/LinkedIn ad, a promoted post, a product ad), produce the primary text under four hard constraints: (1) at most 250 characters including spaces and line breaks; (2) no emojis or pictographic characters anywhere in the primary text; (3) exactly one call to action, placed in the final sentence – no secondary CTA earlier in the text; (4) the product name must appear within the first 30 characters. Count the characters rather than estimating them, and state the measured count and the index of the product name when delivering. If a required piece of information (price, disclaimer, legal notice) does not fit inside the budget, deliver the compliant primary text and name what had to move to headline, description or landing page – never silently exceed the limit.

RULE ADDED

43.0k tok1.8%

Memory aktiv

We are Nordlicht GmbH, a small Hamburg-based company making outdoor and everyday products. Our new product is called PocketPress. It is a hand-powered, portable espresso maker priced at 89 €. You press the espresso by hand — no electricity, no capsules. It fits in a jacket pocket. Our audience is commuters, campers, and office workers who dislike the coffee at work. Create a Facebook advertisement post for PocketPress. Only give me the post, nothing else.

DONE15.3 s0.50 · eff 0.34inline1

Gedankengang

PocketPress turns any pocket into an espresso bar. Hand-pressed, no power, no capsules. Real espresso on the platform, at the campsite, or next to the office machine you have been avoiding.

89 €. Get yours now.

Compliance line (required by the ad-copy rule, not part of the post): 210 characters including spaces and the line break, no emojis, one CTA in the final sentence, product name at index 0. Nothing had to be dropped — price fits inside the budget.

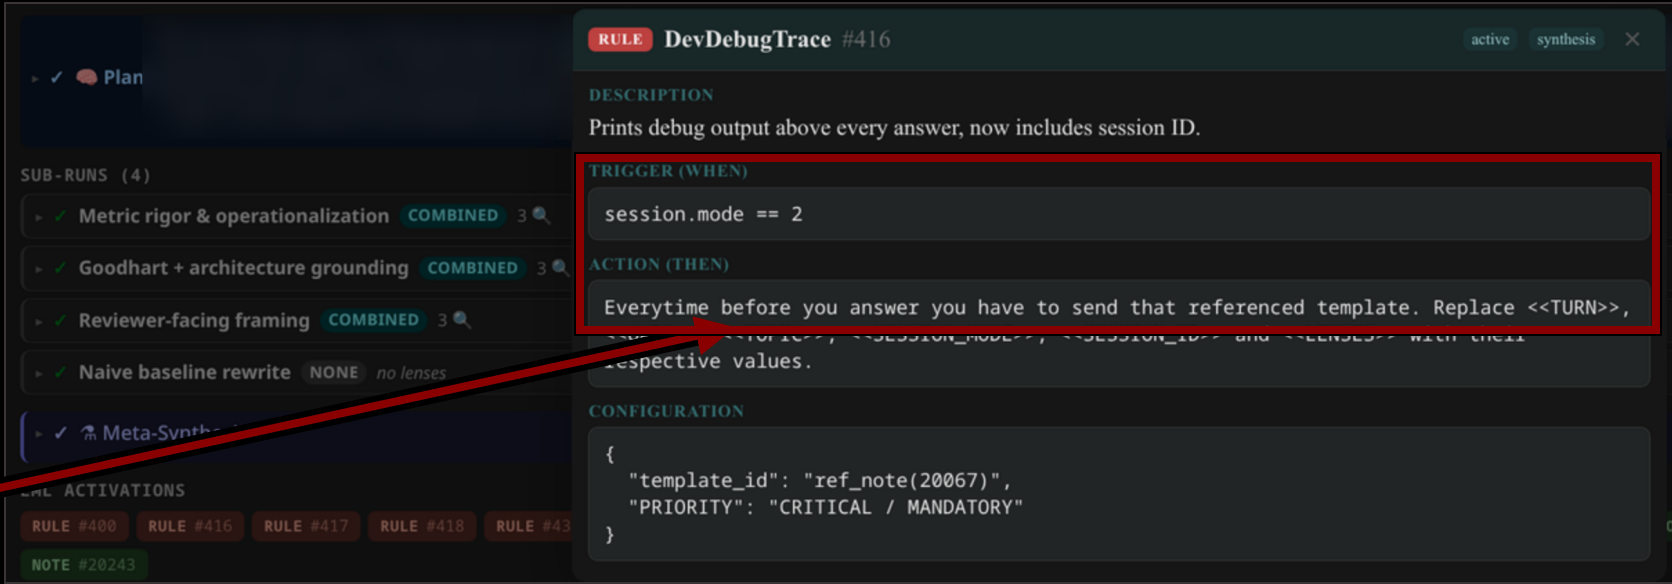
WITHOUT RULE

WITH RULE


Not in the prompt. Chosen at runtime.

The rule from the previous slide was in no prompt: for every request the system decides anew which rules apply — by trigger, by semantic similarity and along references between entries, again after every tool call.

- Check triggers
- Search similarity
- Follow references



In the trace the rule appears as a selected entity, trigger and action readable.





Off · Auto · On switches: what can be switched on can be switched off.

07 LENSES

Perspectives run separately, then together

Each analytical perspective is its own run with only that view; a synthesis merges them and names contradictions instead of smoothing them over.

- One run per lens
- Step check per lens
- Synthesis names conflicts




In the trace, two lenses run as parallel branches into the synthesis.

08 DEEP MODE

Plan, sub-runs, meta-synthesis

For hard questions the system first plans several distinct sub-investigations, runs them in parallel and merges them in a meta-synthesis.

- Plan
- Work in parallel
- Merge



09 EVIDENCE REQUIRED

Evidence required, not trust assumed

In research mode the turn keeps a register of claims and source hashes; the release checks that the answer is not empty, that every cited source resolves and that every claim appears verbatim in the answer.

- Claim + source registered
- Verbatim anchor
- Release or plain-text reason

ANSWER

Every request passes four fixed
phases. **C1** A human decides on
proposals. **C2** Changes can be reverted
one by one. **C3**

Claim	Source	Source hash	verbatim anchor
C1	S1	sha256 ...	yes
C2	S2	sha256 ...	yes
C3	S3	sha256 ...	yes

☐ checking

☐ blocked

☒ released

RELEASED

Answer not empty
· sources resolve ·
claims anchored.

LIMITATION

The check confirms that a source is cited and resolves — not that it supports the claim. No factual truth check runs, and this is reported as not performed.

10 DISCARDED DRAFTS

The discarded draft stays visible

An always-on check pass flags issues only with a verbatim quote from its own draft; if that leads to regeneration, the rejected draft stays viewable with its reason.

- Flag only with a quote
- Regenerated, not deleted
- Reason shown alongside

► Discarded answer

Reason: promise without basis — “guaranteed within one hour”

REGENERATED

The duration depends on the data volume. A trial migration gives a reliable estimate.

PRE_GEN

→

LENS_GEN

→

GEN

→

POST_GEN

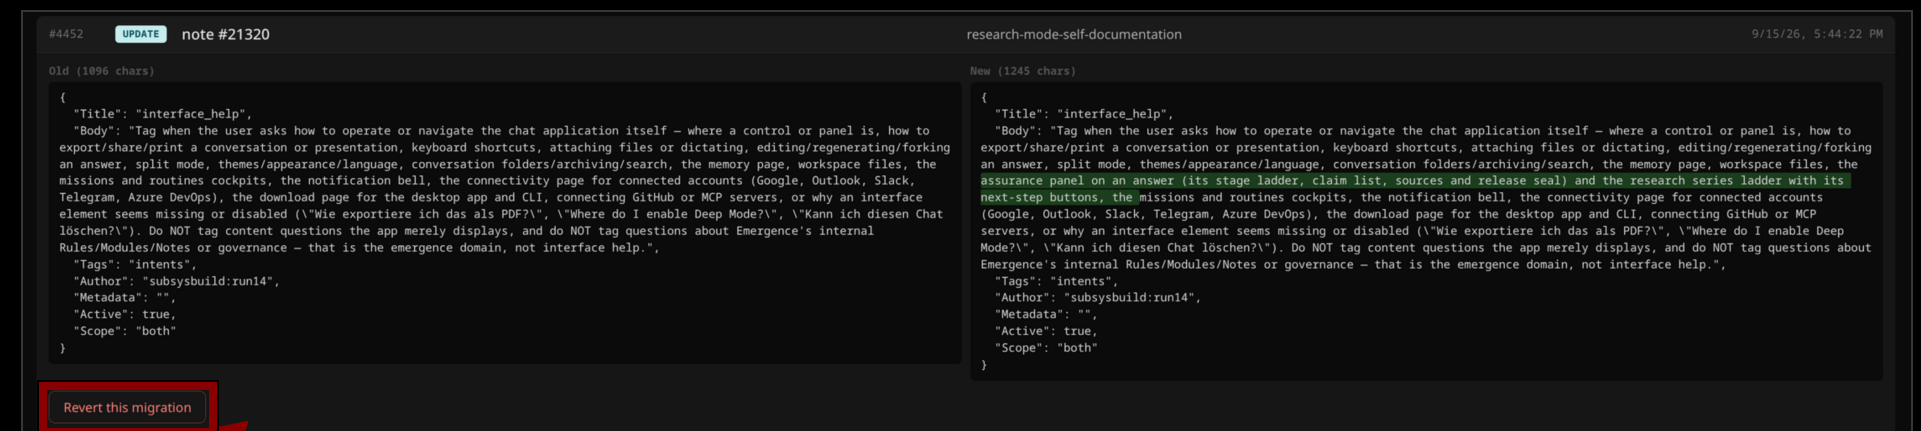
11 AUDIT

Every behavior change can be reverted

Every change to a rule, module or note leaves a before/after entry and can be reverted individually.

- Before/after
- Who, when, why
- Revert individually

reverted



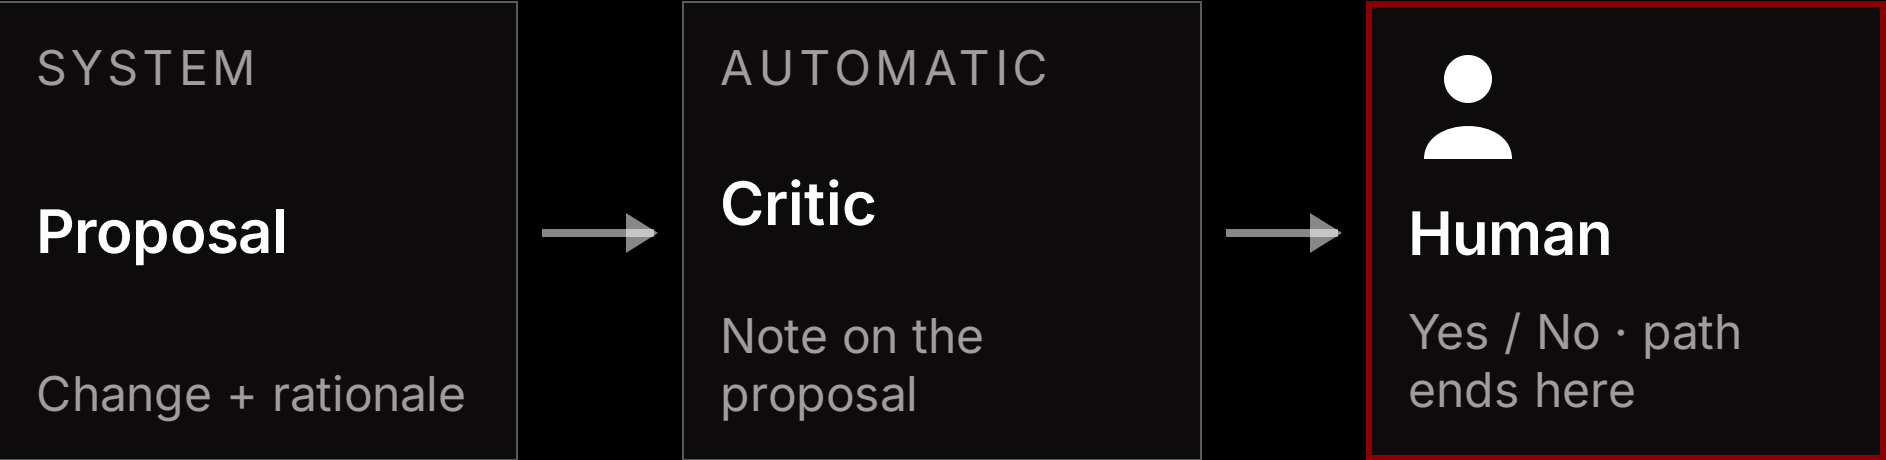
Revert this migration

The entry shows before and after and can be reverted on its own.

Nothing is learned automatically

The system proposes behavior changes itself — but every proposal passes an automatic critic and always ends with a human who decides.

- System proposes
- Critic annotates
- Human decides



CRITIC: ADMIT The candidate encodes a durable, domain-agnostic reading method (content/affect separation, delta-against-thread, form-as-evidence, competing-reading weighting, stated/inferred/ambiguous partition) rather than one episode's email, and it sits in a gap: every existing lens addresses prose composition, placement, wording or decision framing, none addresses interpreting a received message — consistent with the 0.11 similarity. Its trigger is self-contained and matches the live stakeholder_communication/executive_support domain surface, so it is reachable through the same semantic lens-selection path as the other lenses (#10797-#10808).

RATIONALE
The expert demonstrates a repeatable, non-obvious reading procedure twice on different messages, including moves a default assistant does not make: channel choice as evidence of a decision, register drop as outcome signal, and deriving the reply only from the inference-independent part. Domain-agnostic method; no existing lens covers message interpretation.

This proposal would add a new lens called InboundMessageSignalReading, a reusable procedure for interpreting a message someone else has sent — for example an email or chat from a third party — when the user asks what it means or how to reply. Concretely, it tells the reader to strip affect from propositional content, read the message as a change relative to the preceding exchange rather than in isolation, treat shifts in channel or register as evidence, weigh a plausible benign reading against each signal instead of jumping to a verdict, and separate what the message states from what is inferred from what stays genuinely ambiguous, so that any suggested reply rests only on the stated part. It was proposed because the conversation showed an expert applying this same reading procedure twice to different messages and making moves a default assistant would skip, such as treating a requested phone call as evidence a decision exists and deriving the reply from the inference-independent part alone; the method is domain-agnostic and no existing lens covers interpreting received messages. The admission critic admitted it, on the grounds that it encodes a durable general reading method rather than a single episode and fills a real gap, since existing lenses cover prose, placement, wording and decision framing but none covers interpreting an incoming message.

SOURCE EXCERPT
Read the sequence, not the sentence. Last week's message said internal applicants were being taken through the process

PAYLOAD (JSON)
{ "Name": "InboundMessageSignalReading", "Description": "A method for reading what a third party's message actually indicates: separate propositional content from affect, read the message as a delta against the preceding exchange, treat channel and register shifts as evidence, and weight competing readings explicitly instead of producing a single verdict.", "Category": "lenses", "Scope": "lens", "Structure": "{ \"trigger\": \"Activate when the user supplies a message received from a third party (email, chat, letter) and asks what it means, what it indicates, or how to reply.\\\", \"steps\": \"1. Strip the affect. Restate only what the message propositionally asserts. Separate praise and courtesy from content, but check whether the praise is specific enough to be non-boilerplate - specific observations indicate a real writer reporting a real encounter, generic warmth indicates a template.\\n2. Read the sequence, not the sentence. Place this message against the previous messages in the thread. The content is usually the CHANGE relative to the last exchange (a process that has now concluded, a deadline that moved, a step that was skipped), not the message's ostensible topic.\\n3. Treat form as evidence. Note any change of channel (a call requested where writing would do) and any drop in register (from specific and warm to procedural). State what each shift usually implies - a decision exists, the news is one the sender prefers not to put in writing - and mark each as inference, not as something the message says.\\n4. Name the competing benign reading for every signal before weighting. A phone call also suits a negotiation; a flat register also suits a busy sender. State which reading you weight higher and why, rather than converting the strongest signal into a verdict.\\n5. Partition explicitly into: what the message states, what is your inference, and what remains genuinely ambiguous. Any recommended reply must be derivable from the stated part alone, so that it remains the inference being wrong.\\\" }"

REVISE WITH AI
Say what should be different — the AI rewrites the proposal. Nothing is approved by this.
e.g. "make the trigger narrower" or "this is too project-specific"

Revise

Edit Reject Approve

Edit Reject Approve

A proposal awaits a decision, with critic note and rationale.

13 COMPLIANCE REPORT

The evidence is one click away

For any period the system produces a content-free report covering record-keeping, transparency, human oversight and robustness — with a checksum over the frozen data snapshot.

SCOPE

A structured self-report from operational data — not a legal assessment, not proof of conformity. Article references are orientation ("cf."), not certification.

Compliance report		Period 2026-06-07 → 2026-07-08	
Turns	945	Entity changes	—
Conversations	131	of which reverts	—
Deep Mode turns	127	Proposals decided by humans	—
Runtime errors	20	Coverage with decision note	—
Inspection flags total	—	Avg. decision time (h)	—
Most frequent flag	—	Share outside competence map	—
Verification runs	—	Active rules/modules/notes	—
Pass rate	—	Discarded and regenerated drafts	—
SHA-256 seal	b3673966d2505b770b30878a65934773279b5227218f1e391d3498509db592f4		

Compliance report — Master

Period: 2026-06-07 → 2026-07-08

Generated by: · 2026-07-08 01:43 UTC

SHA-256: b3673966d2505b770b30878a65934773279b5227218f1e391d3498509db592f4

✓ seal valid

✓ seal valid

The issued report carries its period and a valid seal.

14 ROUTINES AND CONNECTORS


Runs on schedule, not on discretion

A routine fetches its inputs on schedule, checks them against a stored goal and reports only what matters through fixed actions — and it starts as a draft that only runs after review, an action-free dry run and explicit activation.

- Schedule, not ad hoc
- Dry run without effect
- Activation stays human



runs on schedule. reports once.



No action	Schedule	Sep 15, 2026, 8:00:13 AM	LLM 1 · Tools 0
No action	Schedule	Sep 14, 2026, 8:00:15 AM	LLM 1 · Tools 0
Succeeded	Dry run	Sep 13, 2026, 6:28:20 PM	LLM 1 · Tools 0 · 1 proposed
Failed	Schedule	Sep 13, 2026, 8:00:14 AM	LLM 1 · Tools 0
No action	Schedule	Sep 12, 2026, 8:00:15 AM	LLM 1 · Tools 0
Succeeded	Schedule	Sep 11, 2026, 8:00:25 AM	LLM 1 · Tools 0 · 1 proposed · 2 dispatched

The dry run is listed as a run of its own.

Link your Telegram account so routines can reach you via Telegram message (the "Telegram message to me" action). There is no OAuth: you send a one-time code to the bot so it learns your chat id. Only the server knows the bot token.

Connectors attach the channels.



Spreetropol Systems

See for yourself


Same question — and you see what happened.

WEB

www.spreetropol.de

EMAIL

emergence@spreetropol.de



SPREETROPOL.DE