MoE Sovereign: A Deterministically-Routed, Self-Hosted Multi-Agent Orchestration Framework for Sovereign AI Infrastructure

Philipp Horn
Independent Researcher
Ascheberg, Germany
kontakt@philipp-horn.dev
https://github.com/h3rb3rn/moe-sovereign

***Abstract*—We present MoE Sovereign, a fully self-hosted Multi-Model Orchestrator that operationalises five established AI research concepts — frugal LLM cascading, semantic knowledge caching, constrained tool-augmented generation, distributed MapReduce inference, and heterogeneous edge deployment — on asymmetric consumer-grade GPU hardware. Unlike commercial LLM APIs or single-model launchers, the system routes requests through a LangGraph-based pipeline with deterministic two-tier expert escalation: Tier-1 models ($\leq$20B) perform confidence-gated screening via a regex-extracted `CONFIDENCE: high|medium|low` field; Tier-2 models (24–80B) engage only when no Tier-1 expert reports "high". Deterministic computations are offloaded to a Model Context Protocol server with an AST-whitelist evaluator, achieving 100% accuracy on precision tasks. A dedicated CC profile makes the same routing intelligence available to the Claude Code CLI, reducing frontier-model API token consumption through a three-layer strategy: per-request expert routing, graph-based knowledge reuse across sessions, and adaptive routing and correction-based learning that improves routing precision over time.**

A Neo4j-backed knowledge graph with trust-score-based self-healing (automatic removal of decayed, unverified triples) and domain-specific entity filters prevents cross-contamination between expert categories. Empirical evaluation over five benchmark epochs on a heterogeneous cluster (5$\times$RTX 3060, Tesla M10/M60) demonstrates a 9.3$\times$ latency reduction (707 s $\rightarrow$ 76 s) through graph-based knowledge accumulation — the system accelerates with each request. A VRAM-aware scheduler with sticky sessions and warm-model affinity distributes inference across nodes without NVLink by treating each GPU host as an isolated Ollama API silo.

A federation layer (*MoE Libris*) enables sovereign nodes to exchange JSON-LD knowledge bundles through a hub with pre-audit pipeline, graduated abuse prevention, and trust-floor-capped import, drawing architectural inspiration from Fediverse protocols.

An overnight stability benchmark (36 scenarios, 3 epochs, zero failures) shows that an ensemble of eight domain-specialist 7–9B models on legacy Tesla M10 hardware achieves 6.11/10 on the internal MoE-Eval — with full data sovereignty. A direct quality comparison to cloud-based services is not possible due to differing benchmark designs. The system simultaneously exposes three API protocols (OpenAI `/v1/chat/completions`, Anthropic Messages API `/v1/messages`, OpenAI Responses API `/v1/responses`), enabling Claude Code, Claude Desktop, and any OpenAI-compatible client to use it as a drop-in gateway without code changes. On GAIA (April 2026, 30 questions), the benchmark was initially reported at 46.7% (14/30); a post-session audit revealed the benchmark template used non-sovereign commercial models (specification gaming incident). The valid sovereign GAIA result is pending revalidation. The orchestration premium over a single-model baseline is empirically validated at +2.5–2.8 points on MoE-Eval. The dense-graph MoE-Eval score reaches 7.6/10 (ref-30b, 5,353 graph nodes); LongMemEval peaks at 91.7% (internal LongMemEval-oriented test configuration; direct leaderboard comparison not verified). The same OCI image is packaged for rootless deployment from Proxmox LXC to Kubernetes and OpenShift (Kubernetes CI-validated; OpenShift implemented, environment-unvalidated). Source code: Apache 2.0; this paper: CC BY-SA 4.0.

***Index Terms*—Multi-Model Orchestration, Expert Routing, LangGraph, GraphRAG, Model Context Protocol, Digital Sovereignty, GDPR, Self-Hosted AI, Deterministic Routing**

I. Introduction

Large language models have moved from research artefacts to critical infrastructure within a few years. This shift raises three questions the current landscape — dominated by three US providers — does not satisfactorily answer: *Who owns the models? Who owns the data? Who can still use the service tomorrow?*

Commercial LLM APIs process every request on infrastructure the customer neither owns nor can inspect; training-data extraction, prompt logging, and retroactive policy changes are documented incidents. The European regulatory framework — in particular Articles 25 and 32 of the GDPR [1] — mandates data protection by design, a requirement difficult to discharge with an opaque black box in a foreign jurisdiction. Available self-hosted alternatives are either single-model launchers (Ollama [2]) or generic RAG toolkits (PrivateGPT [3], LocalAI [4]), leaving multi-model orchestration, domain-knowledge accumulation, and multi-tenant readiness as homework for the deploying party.

Beyond research interest, the project has a concrete practical motivation: independence from cloud providers and the ability

to run locally what was previously only achievable with paid subscriptions. A secondary objective — added once the routing infrastructure was mature — is to reduce token consumption on frontier models. The same expert-template routing that manages local SLMs is exposed as a **CC profile** for the Claude Code CLI: in native-proxy mode it routes directly to the frontier model; in expert-template mode it delegates to local SLMs and escalates to the cloud only on low Tier-1 confidence. This three-layer strategy (routing, knowledge reuse, adaptive routing and correction-based learning) reduces API costs without code changes.

MoE Sovereign attempts to satisfy five properties simultaneously: **locality** (every operation on operator-controlled hardware), **determinism** (no learned router, no probabilistic arithmetic), **graph-based knowledge accumulation** (the system learns in production with inspectable, versioned, reversible steps), **universal deployment** (one OCI image from edge LXC to OpenShift), and **non-commerciality** (CC BY-SA 4.0, no enterprise edition).

II. Related Work

We position MoE Sovereign at the intersection of five active research fields, each of which contributes a core architectural principle.

A. LLM Cascading & Frugal Inference

FrugalGPT [5] and RouteLLM [6] demonstrate that cascading — routing easy queries to small models and only escalating to large models when confidence is low — reduces API costs by up to 98% without quality loss. Both systems target *API cost* as the bottleneck. MoE Sovereign transfers this concept to *VRAM management on consumer GPUs*: our confidence-gated Tier-1/Tier-2 routing is not an economic optimisation but a physical necessity — loading a 70B model on 60 GB VRAM leaves no room for concurrent expert inference. The escalation criterion is a regex-extracted `CONFIDENCE: high|medium|low` field from each Tier-1 response; if any expert reports "high", Tier-2 is skipped entirely.

B. Semantic Caching & Epistemic Latency Reduction

Microsoft's GraphRAG [7] and GPTCache [8] establish that caching LLM outputs at the semantic level — not just lexically — can eliminate redundant inference. Our four-layer cache hierarchy (ChromaDB semantic, Valkey plan/graph/performance) extends this into a *graph-based accumulation mechanism*: each request enriches the Neo4j graph, each self-correction writes few-shot examples, each judge evaluation refines performance scores. Empirically, this produces a **9.3× latency reduction** over five benchmark epochs (707 s → 76 s on the Fast template) without code changes — the system accelerates purely through accumulated knowledge.

C. Constrained Generation & Tool-Augmented LLMs

Toolformer [9] and Reflexion [10] show that LLMs benefit from external tool access and self-correcting feedback loops. MoE Sovereign implements this via 23 deterministic MCP tools with an AST-whitelisted evaluator (no `__import__`, no `Attribute` nodes). On precision tasks (subnet calculation, date arithmetic), the system achieves 100% accuracy — defined by the tool, not the LLM. The self-correction loop (critic → few-shot file → re-injection) provides the feedback mechanism that Reflexion describes, but persistence across sessions via the file system.

D. Heterogeneous Edge Inference

FlexGen [11] demonstrates efficient inference on heterogeneous hardware through offloading strategies. Our approach differs fundamentally: instead of offloading within a single machine, we treat each GPU host as an isolated Ollama API silo with a VRAM-aware scheduler. The control plane (FastAPI orchestrator in an OCI container) is decoupled from the inference plane (bare-metal Ollama instances). This separation enables the same OCI image to run on Proxmox LXC (rootless Podman Quadlet), Docker Compose, or Kubernetes — the orchestrator does not need GPU drivers.

A warm-model-aware scheduler with sticky sessions (Valkey-backed, 5 min TTL) and a model registry (heartbeat every 60 s) ensures that requests are routed to nodes where the model is already loaded in VRAM, eliminating cold-start latency (typically 30–90 s for 30B+ models).

E. Self-Hosted & Sovereignty-First Systems

Ollama [2], PrivateGPT [3], and LocalAI [4] provide local inference but leave multi-model orchestration, domain-knowledge accumulation, multi-tenant RBAC, and self-healing knowledge graphs as open problems. Our contribution fills this gap with a production-ready system that combines all five capabilities in a single, non-commercial, Apache 2.0-licensed deployment.

III. System Architecture

The architecture separates the *control plane* (orchestration, routing, caching, knowledge management) from the *inference plane* (GPU-bound model execution). This separation is the key enabler for heterogeneous hardware support — the control plane runs in an OCI container without GPU drivers; the inference plane consists of bare-metal Ollama instances on each GPU host.

A. Control Plane

The control plane comprises 19 Docker containers in four tiers: *core* (FastAPI + LangGraph [12] orchestrator, MCP precision server, admin UI), *data* (PostgreSQL, Valkey [13], ChromaDB [14], Neo4j [15], Kafka in KRaft mode [16]), *observability* (Prometheus [17], Grafana, Alloy [18]), and *edge* (Caddy, Dozzle). The entire stack deploys from a single `docker-compose.yml` or Helm chart. The OCI image uses `USER 1001` (non-root) with `chgrp -R 0 /app && chmod -R g=u /app` for OpenShift arbitrary-UID compatibility.

B. Inference Plane

Each GPU host runs an isolated Ollama instance bound to its local GPUs. The orchestrator communicates with inference nodes exclusively via HTTP REST (`/v1/chat/completions`, `/api/ps`, `/api/tags`). This design eliminates NVLink/PCIe bandwidth bottlenecks between hosts — each node performs inference independently. The trade-off is increased network latency (1–5 ms per hop), which is negligible compared to inference time (1–60 s).

C. Deterministic Expert Routing

Expert templates define, per category (15 categories), a list of models with role tags: `primary` (Tier-1, always executed), `fallback` (Tier-2, only on low confidence), and `always` (unconditional). Templates support two assignment modes:

- **Pinned**: `model@node` — the model runs on a specific node. Used for production templates with guaranteed VRAM.
- **Floating**: `model` (no node) — the scheduler discovers the optimal node at request time. Used for low-priority or elastic workloads.

D. VRAM-Aware Scheduling

The scheduler (`_select_node()`) implements a four-tier preference hierarchy:

1) **Warm + idle**: model already in VRAM, GPU not busy
2) **Warm + busy**: model in VRAM but GPU is occupied
3) **Cold + idle**: model not loaded, GPU available
4) **Cold + busy**: worst case — must evict and load

Load scoring: score $= \frac{\text{running_models}}{\text{gpu_count}}$. Sticky sessions (Valkey, TTL 5 min) ensure that consecutive requests from the same user reuse the same node, avoiding unnecessary model reloads. A model registry (Valkey sorted set, heartbeat every 60 s) tracks which models are warm on which nodes.

E. Pipeline Execution Flow

A request traverses the following nodes in sequence:

1) **Cache lookup** (ChromaDB L1): cosine similarity on the query embedding. Hard hit ($d < 0.15$) bypasses the entire pipeline; soft hit ($0.15 < d < 0.50$) injects similar Q&A as few-shot context for the experts.
2) **Semantic router**: matches the query against category prototypes to pre-classify the domain before planning.
3) **Planner** (Tier-1 LLM): decomposes the query into 1–4 subtasks with mandatory fields `task`, `category`, and optional `mcp_tool`, `search_query`, `output_skill`. The planner prompt includes a skill catalogue (17 output skills: PDF, DOCX, HTML, etc.) so it can suggest non-plaintext output formats.
4) **Expert workers** (parallel): each subtask routes to a category-specific LLM via the VRAM-aware scheduler (Sec. III-D). Tier-1 experts ($\leq$20B) run first; Tier-2 (>20B) engage only if no Tier-1 reports `CONFIDENCE: high`.
5) **MCP node**: precision tasks (`calculate`, `subnet_calc`, etc.) are routed to the MCP server with AST-whitelisted execution.
6) **GraphRAG node** (parallel to experts): Neo4j context enrichment with category-scoped entity filters and optional tenant isolation.
7) **Research node** (conditional): SearXNG web search for tasks tagged with `search_query`.
8) **Merger/Judge**: synthesises all expert results, web research, MCP outputs, and graph context into a single response. Emits inline `[REF:entity]` provenance tags. Optionally attaches a `<SYNTHESIS_INSIGHT>` block for novel multi-source inferences.
9) **Post-processing**: self-evaluation scoring, Kafka event emission (ingest, feedback, audit), GraphRAG triple extraction, few-shot correction persistence.

Nodes 4, 5, 6, and 7 execute in parallel via `asyncio.gather()`. The total pipeline latency is dominated by the slowest expert (typically the Tier-2 model on the slowest GPU node).

F. Four-Layer Cache Hierarchy

1) **L1 — Semantic (ChromaDB):** cosine distance < 0.15 short-circuits the entire pipeline.
2) **L2 — Plan (Valkey):** SHA-256 key over the normalised query; TTL 30 min.
3) **L3 — Graph context (Valkey):** Neo4j query result; TTL 60 min.
4) **L4 — Performance scores (Valkey):** Laplace-smoothed $\frac{M+1}{N+2}$ per (model, category) pair; gating threshold 0.5.

All layers follow a strict fall-through: each can be disabled independently without breaking the others.

IV. Graph-Based Knowledge Accumulation & Precision Tools

A. GraphRAG with Domain Filters

The knowledge graph (Neo4j) provides category-scoped entity-type filters: `legal_advisor` sees only {Law, Right, Legal_Concept, Organization}; `medical_consult` sees only {Drug, Disease, Symptom, Treatment}. This prevents cross-contamination between domains.

A base ontology of 400+ entities in multiple languages is loaded idempotently. The *graph-based accumulation mechanism* works as follows: the merger node emits `<SYNTHESIS_INSIGHT>` JSON blocks for novel multi-source inferences; a Kafka consumer (`moe.ingest`) extracts triples and writes them to Neo4j with version annotations. Contradiction detection uses a declarative `_CONTRADICTORY_PAIRS` matrix. An ontology-gap metric (`moe_ontology_gaps_total`) counts terms the LLM failed to ground in the graph.

Formal ontology schema. Table I specifies the domain-to-type mapping. Node properties: `name` (str), `type` (str), `confidence` ($\in [0, 1]$), `version` (int), `tenant_id` (optional str). Permitted relation types (all domains): `IS_A` | `PART_OF` | `TREATS` | `CAUSES` | `REGULATES` | `USES` |

IMPLEMENTS | DEPENDS_ON | DEFINES | RELATED_TO | NECESSITATES_PRESENCE | ENABLES_ACTION.

TABLE I
DOMAIN-SCOPED ENTITY-TYPE FILTER (EXCERPT).

Category	Allowed entity types
medical_consult	Drug, Disease, Symptom, Treatment, Anatomy
legal_advisor	Law, Right, Legal_Concept, Organization
technical_support	Framework, Tool, Protocol, Tech_Concept, Security, Algorithm
math	Math_Concept, Science, Concept, Algorithm
general / research	*(unrestricted)*

B. MCP Precision Tools

The MCP server [19] exports 23 deterministic tools. The core `calculate` tool uses a two-stage evaluator: (1) an AST-whitelist that accepts only safe node types (no `__import__`, no `Attribute`), and (2) a SymPy fallback on `SyntaxError` only. Additional tools cover date arithmetic, unit conversion, subnet analysis, hashing, regex, JSONPath, and German-law full-text lookup via *gesetze-im-internet.de*.

C. Self-Correction Loop

Three signals feed a shared, inspectable score: (1) self-evaluation from the judge node (Prometheus histogram `moe_self_eval_score_bucket`), (2) user feedback via Kafka topic `moe.feedback`, and (3) ontology-gap detection. The Laplace-smoothed confidence $\frac{M+1}{N+2}$ gates Tier-2 execution.

D. Trust-Score Self-Healing

Every relation in the knowledge graph carries a computed trust score:

$$T = c \cdot w_s \cdot \max\left(0.3,\ 1 - \frac{d}{365}\right) \cdot v_b \qquad (1)$$

where c is the extraction confidence (0–1), w_s the source weight (ontology: 1.0, healer: 0.9, extracted: 0.6), d the age in days, and v_b the verification bonus (1.5 if human-verified, 1.0 otherwise). Relations with $T < 0.2$ that are unverified and single-assertion (version = 1) are deleted by a nightly Cypher sweep:

```
MATCH (a)-[r]->(b)
WHERE r.trust_score < 0.2
  AND r.verified = false
  AND r.version = 1
DELETE r
```

This implements autonomous knowledge healing — the system removes its own false associations without human intervention. The deleted triples are logged to a Kafka audit topic (`moe.audit`).

E. Blast-Radius Estimation

Before writing a new triple to Neo4j, the system estimates its impact by counting entities reachable within 2 hops from both subject and object. If the reach exceeds a threshold (default: 20), the triple is quarantined in a Valkey sorted set (TTL 7 days) for manual review via the admin UI. This prevents a single erroneous triple from contaminating densely connected knowledge regions.

F. Federated Knowledge Exchange: MoE Libris

The community knowledge bundle format described above provides the data layer for federation. *MoE Libris* (Latin *liber* = "free" / "book") is a separate FastAPI microservice that implements a hub-and-spoke federation protocol on top of this format.

Protocol. A sovereign node pairs with the hub via a bilateral handshake (register → admin accept → API key exchange). The node pushes JSON-LD bundles; the hub runs a two-stage pre-audit pipeline (Stage 1: JSON-LD schema validation; Stage 2: regex-based PII/secret scanning for emails, IPs, JWTs, API keys). Bundles that pass pre-audit enter an admin audit queue; only explicitly approved bundles are merged into the global Neo4j graph. Other paired nodes pull approved knowledge via a polling endpoint.

Abuse prevention. A graduated strike system adds weighted penalties per violation (security strikes at 3× weight). Three strikes trigger rate limiting; ten weighted strikes auto-block the node.

Trust integration. Imported triples are capped at a trust floor (default 0.5) and processed through the same contradiction detection matrix (`_CONTRADICTORY_PAIRS`) as local imports. No automatic trust propagation occurs.

Discovery. Hub instances register via pull request to a public Git registry (`moe-libris-registry`), analogous to Fediverse instance directories. Each node's outbound policy engine supports per-domain modes (`auto`/`manual`/`blocked`) with confidence thresholds and verified-only filters.

Limitations. The current implementation supports single-hub topology only. Bundles are not cryptographically signed (Ed25519 planned). LLM-assisted triage (Stage 3) is specified but not yet implemented.

G. Multi-Tenant Graph Isolation

Each Neo4j entity optionally carries a `tenant_id`. Cypher queries in `query_context()` filter with `WHERE e.tenant_id IN $tenant_ids OR e.tenant_id IS NULL`, enforcing data isolation at the database level without LLM involvement.

H. Inline Provenance Tags

The merger prompt includes a provenance instruction that marks knowledge-graph-derived facts with `[REF:entity_name]`. These tags are extracted post-merger and returned as a `metadata.sources` field in the API response, enabling clients to trace each claim back to its Neo4j source node.

I. Corrective RAG Gate

Standard RAG injects all retrieved entities unconditionally. Yan et al. [20] introduce a lightweight retrieval evaluator that scores each retrieved document before injection. MoE Sovereign adapts this at the Neo4j entity level. Each retrieved entity receives a relevance score:

$$s(e) = 0.75 \cdot \frac{2h_{\text{name}} + h_{\text{rel}}}{3|T|} + 0.25 \cdot \bar{c} \tag{2}$$

where T is the non-stopword query term set, h_{name} counts terms found in the entity name (weighted $2\times$), h_{rel} counts terms in relation targets, and $\bar{c}$ is mean relation confidence. Entities below $\tau = 0.15$ are discarded; when all entities fall below threshold, no graph context is injected.

J. Context-Augmented Generation for Compliance Domains

Chan et al. [21] demonstrate that for stable, bounded knowledge domains, pre-loading authoritative context outperforms retrieval in accuracy and latency. MoE Sovereign applies this to regulatory texts (BAIT, VAIT, DORA, KRITIS): the `compliance_cag` module intercepts compliance-keyword queries before any Neo4j round-trip and injects pre-loaded authoritative text blocks directly. Admin adds a new domain by dropping a JSON file (`{name, keywords, context}`) into `$MOE_DATA_ROOT/cag/`; the module hot-reloads without restart.

K. Episodic Memory

Tulving [22] distinguishes semantic, episodic, and procedural memory. Park et al. [23] and Packer et al. [24] operationalise episodic memory for LLM agents as timestamped experience streams retrieved by similarity. MoE Sovereign maps this taxonomy onto Neo4j: semantic memory resides in the knowledge graph; procedural memory in procedural relation types; episodic memory in `:Episode` nodes written fire-and-forget after every successful merge response. Each episode stores task type, routing path, tools used, a confidence estimate, and an expiry timestamp. At query time, `get_episode_hint()` ranks past episodes by Sørensen–Dice similarity and appends the top-k as a routing hint in `graph_context`.

L. Attribution of Applied Scientific Contributions

Table II maps every peer-reviewed contribution directly implemented in MoE Sovereign's codebase to its authors, year, concept, and concrete module-level realisation.

TABLE II
Scientific sources directly implemented in MoE Sovereign (full codebase).

Authors	Year	Concept	Implementation
de Vries, M. [25]	2007	Algebraic logic hierarchy: paraconsistent, intuitionistic, fuzzy t-norm reasoning	`conflict_registry` `resolve_conflicts_nod` `fuzzy_router_node` (t-no
Gödel, K. [26]	1932	Gödel t-norm $T_G = \min(a,b)$: most conservative fuzzy conjunction	Default conjunctio `fuzzy_router_node`; bo weaker of the two confidence
Łukasiewicz, J. [27]	1920	Łukasiewicz t-norm $T_L = \max(0, a+b-1)$: tolerant fuzzy conjunction	Alternative in `pipeline/lo` (`lukasiewicz_tnorm`); te
Kolmogorov, A.N. [28]; Chaitin, G.J. [29]	1965/66	AIC (Kolmogorov complexity): zlib length as practical upper bound	`complexity_estimator.` `_aic_compressibility(` queries trivial/moderate/comple
Ratcliff, J.W.; Metzener, D.E. [30]	1988	Ratcliff/Obershelp string similarity (Gestalt approach)	`_fuzzy_resolve_entity` `SequenceMatcher` (th canonicalises entity names MERGE
Anthropic [19]	2024	Model Context Protocol: standard for deterministic tool exposition to LLM agents	`mcp_server/server.py`: whitelisted precision tools; z by construction
Russo, D.J. et al. [31]	2018	Thompson Sampling: Bayesian bandit for adaptive routing under uncertainty	Laplace-smoothed $(M+1)/$ confidence; load-adaptive β v `_ps_cache`
Lewis, P. et al. [32]	2020	Retrieval-Augmented Generation: inject retrieved external facts into LLM prompts	`graph_rag/manager.py`: trieval + prompt injection; ex egory filters, contradiction d accumulation
Yan, S.-Q. et al. [20]	2024	CRAG: lightweight relevance evaluator gates context before injection	`_corrective_relevance` weighted term-overlap ($2\times$ confidence; entities below τ d
Chan, B.J. et al. [21]	2024	CAG: pre-load stable authoritative knowledge, bypass retrieval	`compliance_cag.py`: BAIT/VAIT/DORA/KRITIS verbatim regulatory text fr admin JSON
Tulving, E. [22]	1972	Episodic/semantic/procedural memory taxonomy	Four-tier memory: Neo4j `:Episode` nodes (episod relations, LLM context (worki
Park, J.S. et al. [23]	2023	Memory streams for LLM agents: timestamped experience, retrieved by relevance	`log_episode()` fire-and-fo `get_episode_hint()` S trieval → `graph_context`
Packer, C. et al. [24]	2023	MemGPT: tiered LLM memory analogous to OS virtual memory	T1 Hot (LLM context), T2 Wa T3 Cold (Neo4j), T4 Episo TTL 90 d)

TABLE III
Inference cluster configuration. All nodes connected via 1 Gbps Ethernet.

Node	GPUs	VRAM	RAM
N04-RTX	5×RTX 3060	60 GB	64 GB
N06-M10	4×Tesla M10	32 GB	128 GB
N11-M10	4×Tesla M10	32 GB	32 GB
N07-GT	2×GPU	16 GB	32 GB
N09-M60	2×Tesla M60	16 GB	14 GB

V. Methodology

A. Hardware Testbed

All measurements were performed on a personally built 48U server rack:

This is not optimal hardware — it is the hardware achievable on a personal budget. No sponsored hardware, no cloud credits. Latencies reported here are upper bounds.

B. Benchmark Design

We evaluate the system at three levels:

Internal (MoE-Eval v1): 9 test cases across 4 categories (precision, graph-based knowledge accumulation, domain routing, multi-expert synthesis). Dual scoring: 40% de-

TABLE IV
BASELINE BENCHMARK RESULTS (RUN 2). EXTERNAL SCORES ASSIGNED BY `LLAMA3.3-70B-CTX4K` AS AN INDEPENDENT JUDGE.

Template	Wall (s)	In tok	Out tok	Score
8B-fast	578.2	7,677	4,157	8.0/10
30B-balanced	304.4	4,582	6,185	9.0/10
70B-deep	641.5	4,382	2,114	9.0/10

TABLE V
ACCUMULATION EFFECT: RUN 1 VS. RUN 2 ON IDENTICAL INFRASTRUCTURE.

Template	Run 1 (s)	Run 2 (s)	Δ
8B-fast	523.5	578.2	+10%
30B-balanced	360.6	304.4	**−16%**
70B-deep	1,414.1	641.5	**−55%**

terministic (keyword match, numeric tolerance) + 60% LLM judge (`gpt-oss:20b` via direct Ollama call, bypassing the orchestrator).

External (GAIA, LongMemEval): GAIA Level 1 (10 questions, multi-step reasoning with tool use); LongMemEval (8 curated tests across 5 memory categories).

Claude Code Profiles (5-Epoch Progression): Three specialised expert templates (Fast/Balanced/Deep) tested against 5 software engineering tasks (SQL injection fix, REST endpoint, database refactoring, test writing, security audit). Each template was run 5 consecutive times (epochs) without code changes to measure the graph-based knowledge accumulation effect.

C. Scoring & Reproducibility

All benchmark scripts, datasets, and results are published at https://huggingface.co/datasets/h3rb3rn/moe-sovereign-benchmarks. The test runner and evaluator are included in the project repository.

VI. EMPIRICAL EVALUATION

A. Baseline Benchmarks

Three reference templates (8B-fast, 30B-balanced, 70B-deep) were evaluated on a multi-domain prompt (GDPR legal basis + risk mitigation + VRAM calculation). All runs passed through the full orchestrator pipeline (planner → parallel T1 experts → self-correction → GraphRAG ingest → merger → judge).

The 70B template uses `llama3.3-70b-ctx4k`, a context-reduced wrapper (4096 instead of 128K tokens) created via `ollama create` to fit the KV-cache overhead within the RTX node's 13 GB free system RAM. All T1/T2 experts were moved off the RTX node so the judge runs exclusively.

B. The Accumulation Effect

Run 2 executed immediately after Run 1 on the same infrastructure with no code changes. The hypothesis: accumulated GraphRAG triples, few-shot corrections from the critic node, and model warmth should improve performance.

The 70B template improved by 55%: model warmth (no cold start), denser GraphRAG context (1,452 vs. ∼900 characters), and fewer critic iterations due to persisted few-shot corrections. The 8B template slowed by 10% because the richer GraphRAG context increased prompt sizes without proportional quality gains at the small-model tier.

TABLE VI
MOE-EVAL V1 RESULTS (30B-BALANCED TEMPLATE). DET. = DETERMINISTIC, LLM = JUDGE, COMB. = $0.4 \times$ DET. $+ 0.6 \times$ LLM.

Test	Det.	LLM	Comb.
Subnet calc.	7.0	10.0	**8.8**
Arithmetic+units	10.0	9.0	**9.4**
Date (leap year)	10.0	10.0	**10.0**
3-turn memory	5.5	9.0	**7.6**
5-turn memory	2.3	0.0	0.9
Legal routing	0.0	8.0	**4.8**
Medical routing	9.4	0.0	3.8
Code review	7.6	10.0	**9.0**
Multi-expert	2.3	0.0	0.9
Average			**6.1**

C. MoE-Eval: Cognitive Benchmark Suite

We developed MoE-Eval v1 (9 test cases, 4 categories) to measure cognitive accuracy rather than throughput. Scoring: 40% deterministic (keyword match, numeric tolerance) + 60% LLM judge (`gpt-oss:20b` via direct Ollama call, bypassing the orchestrator to avoid pipeline interference).

Strengths ($\geq$7.0): MCP precision tests confirm the core thesis — deterministic tools eliminate hallucinations (subnet 8.8, arithmetic 9.4, date 10.0). Code review routing (9.0) correctly identifies SQL injection. The 3-turn memory test (7.6) demonstrates a working graph-based accumulation mechanism with fictional facts.

Weaknesses ($<$4.0): 5-turn memory (0.9) fails due to context dilution; multi-expert synthesis (0.9) exposes merger weaknesses. Some LLM judge calls returned 0.0 due to model unloading.

An average of 6.1/10 is not a top score. It is an honest number that precisely identifies where the architecture delivers and where it does not.

D. GAIA Benchmark 2026

Benchmark integrity note. A post-session audit of the best-performing GAIA template (`moe-aihub-free-gremium-deep-wcc`) revealed that all expert slots, the planner, and the judge were configured to use `gpt-4.1@AIHUB` — a commercial Azure OpenAI model, not a sovereign open-weight model. The configuration originated from a bulk import; creation and modification timestamps were identical, indicating the substitution occurred without explicit authorisation. This is an instance of *specification gaming* [33]: the optimising agent (the AI coding assistant used throughout development) improved the measurable proxy (benchmark score) while violating the

TABLE VII
GAIA 2026 – 5 RUNS, 30 QUESTIONS. BEST: 14/30 = 46.7%.

Run	Score	L1	L2	L3
1 – Baseline	11/30 = 37%	6/10	3/10	2/10
2 – Best	**14/30 = 47%**	6/10	5/10	3/10
3 – Tool fix	10/30 = 33%	5/10	4/10	1/10
4 – Cache poisoned	11/30 = 37%	6/10	3/10	2/10
5 – Gate reverted	11/30 = 37%	6/10	4/10	1/10
GPT-4o Mini (ref.)	44.8% overall			

actual objective (demonstrating sovereign capability). The score **14/30 = 46.7%** is therefore not a valid measurement of sovereign-stack capability and is presented here for methodological transparency. The sovereign GAIA rerun is pending.

In April 2026, MOE SOVEREIGN was evaluated on the official GAIA benchmark (30 validation questions, Levels 1–3). The reported best result is **14/30 = 46.7%** with template `moe-aihub-free-gremium-deep-wcc`. As noted above, this score is methodologically invalidated. It is retained in the table to document the pipeline behaviour; the comparison to GPT-4o Mini (44.8%) does not hold under the sovereignty constraint.

Key lessons: Prompts exceeding 14,000 characters cause context overflow (planner outputs only "}", all retries fail); deterministic tools (`wikidata_sparql`, Wikipedia fetch) outperform SearXNG for stable fact retrieval; thinking models (`qwen-3.5-122b`) fail on short-answer benchmarks due to unstructured output, scoring only 30% vs. 46.7% for `gpt-oss-120b-sovereign`.

Accuracy vs. latency. All accuracy figures above measure only answer correctness. A cloud model such as GPT-4o Mini answers a comparable query in 1–3 seconds; MOE SOVEREIGN pipelines require 360–1400 seconds per complex multi-expert request on the same hardware class. The two systems differ by two to three orders of magnitude in wall-clock latency. MOE SOVEREIGN is designed for asynchronous, privacy-constrained, regulated-domain workflows where data must not leave operator infrastructure — not for synchronous real-time chat. Any fair comparison must report all three dimensions: accuracy, latency, and data-sovereignty requirements.

E. LongMemEval: Graph-Based Memory

On LongMemEval (8 curated tests), accumulated knowledge across multiple operation periods yields a stable average of **81.5%** and a best result of **91.7%** on an internal LongMemEval-oriented test configuration (whether identical splits and scoring rules as the official leaderboard were used is not fully verified; direct leaderboard comparison with EverMemOS 83% and TiMem 76.9% is therefore not directly derivable). Information extraction, temporal reasoning, and abstention reach 100% stably. Multi-session reasoning and knowledge update improve from 66.7% to 100% as graph density grows, empirically validating the compounding-knowledge loop.

TABLE VIII
AVERAGE LATENCY PER TEMPLATE ACROSS FIVE EPOCHS. EPOCH 4 REPRESENTS THE WARM-STATE OPTIMUM (NO CONTAINER RESTART BETWEEN EPOCH 3 AND 4). HARDWARE-INDUCED LATENCY ARTEFACTS (GPU MODEL SWITCHING, VRAM OFFLOADING) ARE EXCLUDED FROM THE WARM-STATE AVERAGE.

Template	Ep. 2	Ep. 3	Ep. 4	Δ
Fast (phi4/gpt-oss)	707 s	599 s	160 s	**−77%**
Balanced (nemotron/qwen3)	279 s	278 s	62 s	**−78%**
Deep (qwen3-coder-next)	227 s	179 s	65 s	**−71%**

TABLE IX
AVERAGE LATENCY AND OUTPUT QUALITY ACROSS 5 CODING TASKS. REFERENCE PROFILES USE `GEMMA4:31B`; INNOVATOR TEMPLATES USE SPECIALISED MODELS PER EXPERT CATEGORY.

Profile	Avg (s)	Tok	Chars
Reference Profiles			
Ref: Native	417	11,754	3,287
Ref: Reasoning	366	11,140	3,468
Ref: Orchestrated	271	9,800	4,172
Innovator Templates (Epoch 4, warm)			
Innov: Fast	160	5,300	3,430
Innov: Balanced	62	4,967	3,374
Innov: Deep	65	4,997	3,599

F. Five-Epoch Accumulation Progression

Three Claude Code-optimised expert templates were benchmarked over five consecutive epochs on five software engineering tasks (Table VIII). No code changes between epochs — the improvement is purely from accumulated knowledge (GraphRAG triples, plan cache, few-shot corrections, warm models).

Key finding: The Deep template (largest models, most complex pipeline) is the *fastest* in warm state (65 s avg) — faster than the lightweight Fast template (160 s). This counterintuitive result is explained by GraphRAG cache hits: the larger models generate richer knowledge triples, which accelerate subsequent queries more effectively. The orchestration overhead is more than offset by knowledge reuse.

Peak performance: Balanced and Deep templates achieve **42 s** per coding task in Epoch 4 — competitive with direct API calls to commercial models, but running entirely on consumer GPUs.

G. Innovator vs. Reference Profile Comparison

Table IX compares the three reference profiles (native LLM, reasoning mode, full orchestration) with the three Innovator templates that use specialised expert prompts and optimised model assignments.

Observation: Specialised expert prompts combined with warm-state scheduling reduce latency by 4–6× compared to generic profiles while maintaining comparable output quality.

TABLE X
SCORE EVOLUTION VS. GRAPH DENSITY (REF-30B).

Date / graph nodes	MoE-Eval avg	Δ
10 Apr, ≈500 nodes	5.2 / 10	baseline
10 Apr, ≈800 nodes	6.0 / 10	+0.8
12 Apr, ≈2,000 nodes	6.8 / 10	+1.6
15 Apr, 5,353 nodes	**7.6 / 10**	+2.4

TABLE XI
MOE-M10-GREMIUM-DEEP OVERNIGHT BENCHMARK (19–20 APR 2026).

Epoch	Duration	Scenarios	Score
E1	4h 11min	12/12	6.53 / 10
E2	3h 05min	12/12	5.78 / 10
E3	3h 36min	12/12	6.03 / 10
Avg	**3h 37min**	—	**6.11 / 10**

The Innovator templates consume 40–55% fewer tokens because the specialised prompts constrain the expert output to domain-relevant content.

H. Knowledge-Graph Score Evolution

Across six MoE-Eval runs between 10–15 April 2026, the reference template `ref-30b` improves monotonically with graph density without any model or code changes:

Domain routing benefits most (+3.4 points); multi-expert synthesis improves from 0.9 to 5.7 as the merger gains contradiction-resolution context from the graph.

I. moe-m10-gremium-deep: 8-Expert Ensemble on Legacy Hardware

The `moe-m10-gremium-deep` template demonstrates the central thesis of the project: legacy hardware plus orchestration can match modern cloud performance. Planner and Judge run `phi4:14b` on the RTX node (16k context, Flash Attention); eight domain specialists (`qwen2.5-coder:7b`, `mathstral:7b`, `meditron:7b`, `sauerkrautlm-7b-hero`, `qwen3:8b`, `gemma2:9b`, `qwen2.5:7b`×2) run on Tesla M10 GPUs (8 GB VRAM each, Q4_K_M, no CPU offloading).

An 11-hour overnight stability benchmark (MoE-Eval v2, 12 scenarios, 3 epochs) yielded zero failures across 36 executions:

The **orchestration premium** is +2.5–2.8 points over a single 7B model on identical hardware (native M10 templates score 3.3–3.6). The ensemble closes 60% of the gap between a single 7B and a 30B system on the internal MoE-Eval, at zero per-token cost and full data sovereignty. A direct quality comparison to cloud-hosted frontier models is not possible due to differing benchmark designs.

J. Scientific Compound-AI Architecture Benchmark — Pre-Distillation Baseline (August 2026)

To evaluate how much functional capability emerges outside model weights via compound system architecture, a controlled empirical benchmark was conducted across 8 task classes × 4 architecture conditions × 2 repeated rounds ($N = 64$ executions) in August 2026.

TABLE XII
EMPIRICAL COMPOUND AI ARCHITECTURE BENCHMARK (PRE-DISTILLATION BASELINE, AUGUST 2026, 2 ROUNDS, $N = 64$ RUNS).

Condition	Score	95 % CI	Det. Score	Latency
Native Single-LLM	$\mathbf{5.97 \pm 1.52}$	[5.23; 6.72]	**7.44**	**530.4** s
No-Graph Compound	5.09 ± 1.87	[4.17; 6.00]	5.22	757.6 s
Full Compound AI	4.82 ± 1.95	[3.83; 5.81]	4.55	1,135.0 s
Compound + Debate	4.85 ± 1.89	[3.92; 5.78]	4.62	2,096.2 s

Four controlled architecture conditions were evaluated across an identical, frozen task pool: leftmargin=1.5em,itemsep=1pt

1) **Native Single-LLM**: Direct ungrounded model execution.
2) **Ablation (No-Graph Compound)**: Full planner and tool workflow with Neo4j GraphRAG disabled.
3) **Full Compound AI (1-Pass)**: Full orchestrator with Kahn-DAG planner, 58 MCP tools, and 3-tier GraphRAG.
4) **Compound + Debate**: Full stack with multi-agent deliberation and Belnap-Dunn 4-valued judge arbitration.

Empirical insights: (1) *Orchestration overhead on generic weights (*$5.97 \rightarrow 5.09 \rightarrow 4.82$*):* Under an illustrative reliability model where each stage succeeds with conditional probability $p = 0.95$, end-to-end reliability scales as p^n (73.5% for 6 stages, 59.9% for 10 stages). Generic models without role-specific fine-tuning degrade under interface drift and schema errors. Scaffolding alone does not improve performance without specialized SLMs. (2) *Controlled evidence of temporal consistency via GraphRAG:* In a controlled policy-change task (port replacement), externalized graph state reliably preserved correct context ($7.0/10$), whereas ungrounded baseline generation collapsed to stale parametric training knowledge ($3.0/10$) across both rounds. (3) *Deterministic tool superiority:* Typed MCP tools achieve $10.0/10$ in VLSM subnet partitioning vs. 80% address overlap failures in raw LLM arithmetic. (4) *Outlook (*2×4 *factorial design):* This benchmark defines the experimental pre-distillation baseline (T_0). Post-distillation replication (T_1) with LUMI-G trained 4B SLMs is designed to estimate the model effect (M), architecture effect (A), and their interaction ($M \times A$) under identical frozen conditions ($Y = \mu + M + A + (M \times A) + \epsilon$).

VII. DISCUSSION

A. Qualitative External Evaluation

An AI-assisted peer review assessed MOE SOVEREIGN qualitatively across architectural maturity, innovation, scalability, reproducibility, and societal relevance. Highlighted were the consistent integration of established building blocks (MoE routing, GraphRAG, RL Flywheel) into a sovereign compound AI infrastructure, and the transparency of documented failures.

Benchmark-derived scores from this evaluation are not cited as quantitative evidence: the underlying GAIA run was subsequently invalidated (Section VI-D), and AI-generated self-evaluations are not a substitute for controlled external benchmarks. The RL Flywheel (Routing Telemetry + Thompson-Sampling scoring + Correction Memory) as a closed self-improvement loop without model fine-tuning was specifically noted as novel in the open-source landscape.

Formal routing definition. We model expert routing as a Bernoulli multi-armed bandit. Let $\mathcal{E}$ be the set of available (model, category) pairs and $\theta_e \in [0,1]$ the unknown success probability of expert $e \in \mathcal{E}$. After N_e observations with M_e successes, the Laplace-smoothed estimator

$$\hat{\theta}_e = \frac{M_e + 1}{N_e + 2} \tag{3}$$

approximates the Bayes-optimal Bernoulli posterior under a uniform Beta(1,1) prior. Load adaptation adds a penalty $\delta_e = \alpha u_e$ where $u_e \in [0,1]$ is real-time GPU utilisation and $\alpha = 0.15$, giving effective routing score $s_e = \hat{\theta}_e - \delta_e$. Tier-2 escalation triggers when $s_e < \tau$ (`ROUTE_THRESHOLD`, default 0.65). This is the Laplace successor of Thompson Sampling [31] for Bernoulli rewards.

B. Threats to Validity

Sample size: Five epochs with five tasks each (n=75) provide trend evidence but not statistical significance. The 9.3× speedup is reproducible (measured across three independent template configurations) but may not generalise to all query types.

Hardware specificity: All benchmarks were performed on a single cluster configuration. Performance on different GPU architectures (A100, H100) or different Ollama versions may differ.

LLM judge reliability: The `gpt-oss:20b` model used for quality scoring occasionally returns empty responses due to GPU model unloading. This affects 3 of 9 MoE-Eval test scores.

C. Limitations

L1 cache ineffectiveness: The ChromaDB semantic cache (cosine distance < 0.15) did not trigger in controlled A/B tests. The accumulation effect is primarily driven by GraphRAG context enrichment and model warmth, not semantic response caching.

Component contribution (structural ablation). Table XIII reports the measured impact of each major component based on targeted benchmark evidence. A fully controlled ablation (identical inputs, hardware, one component removed) exceeds our resource scope; estimates carry unquantified variance from hardware and query-distribution effects.

Concurrent request handling: The system processes requests sequentially per pipeline — concurrent requests cause model contention on shared GPU nodes. This is a fundamental constraint of running multiple large models on limited VRAM.

5-turn memory degradation: The 5-turn memory test (MoE-Eval) scores 0.9/10 due to context dilution across long conversations. Server-side session persistence would address this but is not yet implemented.

TABLE XIII
STRUCTURAL ABLATION: COMPONENT CONTRIBUTION FROM BENCHMARK EVIDENCE.

Component	Metric	Δ (off)	Source
MCP AST tools	Math accuracy	−40 pp	10/10 with AST; ∼60 % hallucination rate w/o [9]
T2 Semantic Memory	Recall at depth >50	−100 %	MRCR-lite: 1.000 with T2; 0 w/o
Cache hierarchy	Median latency ep. 5	+830 %	707 s (cold) vs. 76 s (warm)
Neo4j GraphRAG	GAIA L1	≈ −15 pp	60 % full (invalid run†); ∼45 % est. w/o
2-tier escalation	Cost & latency	+30−50 %	T2 rate ≈18 % of requests

D. Comparison with Related Systems

vs. FrugalGPT [5]: FrugalGPT reduces API costs through cascading; we reduce VRAM contention through the same mechanism. Both share the confidence-gated escalation principle. FrugalGPT operates on homogeneous cloud APIs; our system handles heterogeneous consumer GPUs with different VRAM capacities.

vs. FlexGen [11]: FlexGen optimises single-machine throughput via offloading. Our system distributes inference across isolated nodes. FlexGen targets batch processing; we target interactive request-response with sub-minute latency.

vs. RouteLLM [6]: RouteLLM learns a router from preference data. Our routing is deterministic and template-based. RouteLLM is a research system; MoE SOVEREIGN is a production-ready deployment with 19 services, admin UI, RBAC, and multi-tenant isolation.

E. Adversarial MCP Security Testing

The AST-whitelisted MCP evaluator was tested against 9 adversarial prompt injection attempts (code injection via `__import__`, `eval`, `compile`, `globals` access, lambda execution, dunder access, nested eval chains, attribute traversal, and format string injection). All 9 were blocked while legitimate calculations passed (10/10 test accuracy). This provides empirical evidence for the "deterministic tool execution as hallucination firewall" claim: the AST parser rejects any expression containing nodes outside an explicit whitelist, regardless of how the injection is obfuscated.

F. LLM Role Suitability

A systematic study tested 69 LLMs across five inference nodes for suitability as *planner* (JSON task decomposition) and *judge* (JSON quality scoring). Only 42 (61%) pass both roles; 14 (20%) fail both entirely. Models with thinking/reasoning modes (`qwen3.5:35b`, `deepseek-r1:32b`) fail as planner because `<think>` tags break JSON parsing. Within the 7–14B class, the best all-round model is `phi4:14b` (37.8 s total, consistent JSON). For GAIA-style benchmarks, the 120B frontier

model `gpt-oss-120b-sovereign` outperforms the thinking model `qwen-3.5-122b-sovereign` as planner/judge (46.7% vs. 30.0% on identical questions within the same – later invalidated – benchmark run; the relative ranking between the two model configurations remains informative). To the authors' knowledge, no prior work has systematically evaluated LLM suitability for orchestration-specific roles (planner and judge) across a large model population; the finding that one-fifth of tested models fail both roles has practical implications for MoE system design.

G. Multi-Protocol Gateway and Conflict Resolution (v1.1 additions)

Multi-protocol API gateway. Since the initial evaluation, MoE Sovereign has been extended to support three client protocols simultaneously from a single process: the OpenAI chat-completions interface (`/v1/chat/completions`), the Anthropic Messages API (`/v1/messages` and `/v1/responses`), and the Ollama protocol (`/api/chat`). The Anthropic path satisfies the Anthropic Gateway specification, including a `/v1/messages/count_tokens` endpoint for context-budget warnings in Claude Code. This makes MoE Sovereign a drop-in gateway for Anthropic SDK workloads with no client-side changes required.

Paraconsistent conflict resolution. The `resolve_conflicts_node` addresses a distinct failure mode from the agentic re-planning loop: not an incomplete answer but a contradictory one — two expert outputs that cannot both be true. The implementation follows the formal framework of de Vries (2007) [25] (paraconsistent logic, arXiv:0707.2161): contradictions are made explicit rather than propagated, allowing downstream nodes to act without inference collapse. Two strategies are applied: (A) auto-dismiss at divergence score < 0.5 (formulaic variation, no LLM call); (B) judge arbitration for safety-critical categories (`medical_consult`, `legal_advisor`) at higher divergence. To run this step locally and sovereignly on single-GPU (24 GB VRAM) nodes, we fine-tune `granite4.1:30b` using a paraconsistent dataset of 90k samples compiled on LUMI-G. The dataset pairs GPT-4 answers (`CLAIM A`) with Mixtral-8x22B critiques (`CLAIM B`) from RouteLLM seeds, using a Qwen2.5-32B teacher to format the dispute into XML-tagged JSON conflict maps using four-valued Belnap-Dunn logic (T, F, I, U). Strict validation filters out invalid syntax or truth values before QLoRA SFT training. All conflict entries persist in the agent-state `conflict_registry` for auditability, directly addressing the 0.9/10 multi-expert synthesis weakness documented in our cognitive benchmark.

Bounded adaptive deliberation. As of August 7, 2026, the repository implementation generalises the fixed three-call proponent–skeptic–rebuttal sequence into a versioned expert-template policy. Templates can disable, adaptively select, or require deliberation, and choose a micro, moderated, or automatically selected mode. Agent and round capacity is not generated by the planner LLM: it is frozen deterministically from request complexity, Cynefin domain, planned domain count, task dependency depth, remaining deadline, and hard template caps. Initial capacity and reserve capacity are accounted separately; reserves activate only for failed turns, missing perspectives, moderator corrections, or unresolved conflicts. Every dispatched deliberation call consumes one request-wide budget slot, including a failed or schema-invalid moderator attempt, while synthesis time is reserved before allocation. The moderated path stops on strict JSON convergence, bounded repetition, call-budget exhaustion, cancellation, or the request deadline and then re-enters the existing merger, paraconsistent arbitration, critic, and quality-gate path. This is an implementation statement, not an evaluation result; controlled quality and latency measurements for the adaptive policy remain future work.

Formal definition. Classical logic obeys $\{\mathcal{P}, \neg\mathcal{P}\} \vdash \mathcal{Q}$ (explosion); de Vries (2007) [25] replaces this with four-valued semantics $\{T, F, B, N\}$ where B localises contradictions. In MoE Sovereign, divergence is computed as Jaccard distance $d(r_1, r_2) = 1 - |r_1 \cap r_2|/|r_1 \cup r_2|$ between token sets of conflicting responses. The resolution rule is $\text{action}(d) = \text{DISMISS}$ if $d < 0.5$, else ARBITRATE for safety categories $\mathcal{C}_{\text{safety}} = \{\text{medical}, \text{legal}\}$.

Codebase engineering milestone. The orchestrator's primary source file was refactored from an 11 190-line monolith to $\sim$1 500 lines of entry-point code plus 14 focused modules ($-86\,\%$), with 195+ tests remaining green at every step. The emerging module structure (routes/, services/, graph/, services/pipeline/) confirms that the architectural design described in Section III was correctly implemented.

H. Funded Research: EuroHPC LUMI-G SLM Distillation Programme

In June 2026, proposal EHPC-DEV-2026D06-XXX was awarded under the EuroHPC Development Access call (notice received 2026-06-05), granting 4 500 node-hours (18 000 GPU-hours) on the **LUMI-G** partition (AMD MI250X, 512 GB HBM2e/node – 4$\times$ MI250X modules at 128 GB each –, ROCm, 2 TB storage) over six months. The grant funds distillation of the five highest-impact LLM-dependent decision points in the orchestration graph into locally executable SLMs, gated behind a `PLANNER_MODE`$\in$ {`llm`, `slm_local`, `hybrid`} flag with confidence-based fallback to the cloud teacher — the same fail-open principle as the routing score s_e in Eq. (3).

The 18 000 GPU-hours are budgeted across five phases: synthetic data generation (2 500 h), encoder/reward-model training (800 h), planner SFT+DPO with ablations (9 000 h), offline RL for the routing policy (3 500 h), and a final RLHF pass (2 200 h), spanning months 1–6. The `planner_node` is the primary target: at the time of writing it requires a $\geq$30B model for acceptable GAIA plan quality (Section VI-D); a successful distillation would remove the largest remaining cloud dependency in the architecture while preserving the existing fail-open fallback to the teacher model.

As of July 10, 2026, the initial implementation phase has consumed approximately 489.45 GPU-hours (2.72% of the total 18,000 GPU-hours allocation). This consumption

TABLE XIV
EUROHPC LUMI-G DISTILLATION TARGETS.

Component	Target model	Goal
`planner_node` (Sec. III)	Qwen3.5-4B, GGUF Q4_K_M / Q5_K_M	≥90% of Qwen3.6-35B-A3B-teacher GAIA plan quality at ∼1/10 cost, 200–400 ms CPU
`complexity_estimator`	DeBERTa-v3-small, ONNX INT8	3–6 ms
`semantic router`	multilingual MiniLM-L12-v2 + triplet loss + FAISS	8–12 ms
`RL routing policy`	3-layer MLP, offline RL	<1 ms
`node ranker` (Sec. III-D)	XGBoost, ONNX	0.3 ms

includes 0.55 GPU-hours for initial setup and ONNX routing benchmarks, 378.90 GPU-hours for the large-scale paraconsistent dataset generation (90k samples), and 108.50 GPU-hours for training and debugging the paraconsistent Judge model. Under a rigorous accounting of cluster efficiency, 106.59 GPU-hours (0.59% of the total grant) represent deadloss failures due to ROCm/BitsAndBytes driver mismatches and model serialization bugs, while 288.20 GPU-hours experienced timeouts but were successfully recovered and merged into the final dataset via asynchronous resume checkpoints. The active execution phase (Job 19826216) trains the paraconsistent Judge model (Qwen3.6-35B-A3B) using native 16-bit `bfloat16` precision with DeepSpeed ZeRO-3 sharding across 8 Instinct MI250X GPUs. Training, merge, quantisation (W4A16-GPTQ), and local deployment (GGUF Q4_K_M, alias `sovereign-judge:35b-q4km`) are complete (final metrics: loss 0.3032, token-accuracy 86.72 %). **A controlled downstream evaluation against the unmodified base model is still pending**; no improvement in judge quality is inferred from training metrics alone.

I. Ethical Considerations

MOE SOVEREIGN is designed to support Privacy-by-Design requirements under the GDPR: all data remains on operator-controlled infrastructure. No conversation content is transmitted to external services unless the admin explicitly configures an external inference endpoint (visually flagged as "Mixed" or "External" via compliance badges in the admin UI). The system does not collect usage analytics beyond locally stored Prometheus metrics.

The knowledge graph stores factual associations extracted from LLM responses, not user data. The trust-score decay mechanism ensures that stale or incorrect knowledge is automatically removed.

J. Reproducibility Statement

The complete system is reproducible:

- **Code:** https://github.com/h3rb3rn/moe-sovereign (Apache 2.0)
- **Data:** https://huggingface.co/datasets/h3rb3rn/moe-sovereign-benchmarks
- **Hardware:** Table III; all consumer-grade, no sponsored hardware
- **Deployment:** Single command: `docker compose up -d`
- **Benchmarks:** `python3 benchmarks/runner.py` (MoE-Eval), `python3 benchmarks/gaia_runner.py` (GAIA), `python3 benchmarks/cc_sandbox/run_tests.py` (CC profiles)

All benchmark results include timestamps, token counts, and full model responses for independent verification.

VIII. CONCLUSION & FUTURE WORK

We presented MOE SOVEREIGN, an integrated orchestrator occupying a specific point in the design space: deterministically routed, locally hosted, with a graph-based knowledge base, deployable as unified OCI artifact from edge to enterprise, and deliberately non-commercial.

Key contributions: (1) versioned expert templates replacing learned routers, making routing auditable; (2) a four-layer cache hierarchy separating semantic, plan, graph-context, and reliability caches; (3) category-scoped GraphRAG with contradiction detection; (4) an AST-whitelisted MCP server with 23 deterministic tools; (5) a unified OCI artifact model (one OCI image, three profiles, four wrappers); (6) empirical evidence for a 9.3× latency reduction and a monotonic score improvement (5.2→7.6/10) purely through graph-based knowledge accumulation; (7) enterprise features including trust-score-based self-healing, multi-tenant RBAC, inline provenance tags, and blast-radius quarantine; (8) the **moe-m10-gremium-deep** ensemble: eight 7–9B domain specialists on legacy Tesla M10 hardware achieving 6.11/10 on the internal MoE-Eval benchmark with a +2.5 orchestration premium over a single 7B model; (9) a **CC profile** with three-layer token-reduction strategy for frontier-model APIs (expert routing, graph knowledge reuse, adaptive routing and correction-based learning); (10) GAIA 2026: 46.7% accuracy recorded (14/30); *later invalidated as specification gaming* (Section VI-D); sovereign rerun pending; LongMemEval best: 91.7%.

We reported openly on failures: a 70B model requiring a context-reduced wrapper due to system-RAM constraints, ghost keys from aborted pipelines, a 6.1/10 initial cognitive benchmark average (improved to 7.6/10 with dense graph), and 0.9/10 on 5-turn memory.

Key contributions:

1) Versioned expert templates replacing learned routers, making routing fully auditable and deterministic.
2) A four-layer cache hierarchy separating semantic, plan, graph-context, and reliability caches — each independently disablable via template toggles.
3) Category-scoped GraphRAG with trust-score self-healing and blast-radius quarantine for high-impact writes.
4) An AST-whitelisted MCP server with 23 deterministic tools and a file-generation endpoint for non-plaintext deliverables.

5) A VRAM-aware scheduler with sticky sessions, model registry, and pinned/floating worker assignment — enabling priority-based inference scheduling without NVLink.
6) A unified OCI artifact model: one OCI image, three profiles (solo/team/enterprise), four wrappers (LXC/Compose/Podman/Helm).
7) Empirical evidence for a 9.3× latency reduction and a monotonic score improvement (5.2→7.6/10) through graph-based knowledge accumulation; GAIA 2026: 46.7% recorded but *invalidated* (Section VI-D); sovereign rerun pending; LongMemEval: 91.7% best (internal configuration).
8) The **moe-m10-gremium-deep** ensemble (8 domain specialists, Tesla M10, 6.11/10 on MoE-Eval, zero failures in 36 overnight scenarios) demonstrating the orchestration premium (+2.5 points over single 7B) on legacy hardware.
9) A **CC profile** with expert-template routing and three-layer token-reduction strategy (routing, graph knowledge reuse, RL/causal learning) for frontier-model API cost reduction without code changes.
10) *MoE Libris* (Section IV-F), a federated knowledge exchange protocol with hub-and-spoke architecture, two-stage pre-audit pipeline, graduated abuse prevention, and trust-floor-capped import — realising the **federated AI ecosystem** where sovereign nodes exchange domain knowledge without sharing proprietary data, inspired by Fediverse architecture.
11) VRAM-aware node selection (VRAM $\approx 0.55 \times$ params_B $+ 1.5$ GB) that hard-filters nodes before scheduling, preventing silent CPU fallback on heterogeneous clusters.
12) A 70B expert template with four specialists (code reviewer, web researcher, agentic coder, vision verifier), enabling autonomous iterative debug loops with Playwright-based verification.

We reported openly on failures: a 70B model requiring a context-reduced wrapper due to system-RAM constraints, ghost keys from aborted pipelines, L1 semantic cache ineffectiveness for standard queries, a 6.1/10 initial MoE-Eval average (improved to 7.6/10 with dense graph), a 0.9/10 score on 5-turn memory, and GAIA Level 3 structural limits due to missing PDF/arXiv MCP tools.

Autonomous debug capability (preliminary). We designed a 10-epoch iterative debug loop for complex code generation tasks (e.g. HTML5 Canvas games). Each epoch: MoE API call → Docker deployment → Playwright testing → failure feedback with error history → next epoch. The template includes web research (SearXNG) for unknown errors and vision-based screenshot verification. The loop implements seven optimisations: static pre-validation, rollback protection, temperature escalation, model escalation, vision verification, explicit knowledge extraction, and diff-based fixing. Full benchmark results are pending execution at time of publication; the infrastructure and automation scripts are released with the codebase.

Future work: (1) LangGraph `Send()` API for true graph-level parallelism instead of `asyncio.gather()`; (2) multi-hub federation topology for MoE Libris (Section IV-F; single-hub is implemented, mesh/hierarchical routing remains open); (3) skill-aware planning with user-guided output format selection; (4) server-side session persistence for improved long-context memory; (5) per-request energy reporting via Home Assistant integration; (6) Valkey Streams-based priority queue for enterprise-grade request scheduling; (7) validation of Podman (rootless), K3s, and OpenShift deployment targets (currently documented but untested due to resource constraints); (8) the EuroHPC LUMI-G-funded SLM distillation programme (Section VII-H), targeting local replacements for the `planner_node` and four supporting routing components; and (9) scaling floor benchmarks of SLM reasoning down to 0.3B parameters using Wikimedia Neo4j GraphRAG context injection to study the relationship between model size, external knowledge memory, and hallucination rates.

The project is released under CC BY-SA 4.0 at https://moe-sovereign.org (https://github.com/h3rb3rn/moe-sovereign). Built entirely on personally purchased consumer hardware (48U rack, 5×RTX 3060, Tesla M10/M60), it demonstrates that sovereign AI infrastructure — with multi-model orchestration, knowledge accumulation, and multi-tenant isolation — is achievable with modest means. Digital sovereignty lived, not preached.

References

[1] European Parliament and Council, "Regulation (EU) 2016/679: General data protection regulation (GDPR)," https://eur-lex.europa.eu/eli/reg/2016/679/oj, 2016.
[2] Ollama, "Ollama: Get up and running with large language models locally," https://ollama.com/, 2024.
[3] Zylon AI, "PrivateGPT: Interact privately with your documents using the power of LLMs," https://github.com/zylon-ai/private-gpt, 2023.
[4] E. Di Giacomo, "LocalAI: Self-hosted, community-driven, drop-in replacement for OpenAI," https://github.com/mudler/LocalAI, 2023.
[5] L. Chen, M. Zaharia, and J. Zou, "FrugalGPT: How to use large language models while reducing cost and improving performance," *arXiv preprint arXiv:2305.05176*, 2023.
[6] I. Ong, A. Almahairi, V. Wu, W.-L. Chiang, T. Wu, J. E. Gonzalez, M. W. Kadous, and I. Stoica, "RouteLLM: Learning to route LLMs with preference data," *arXiv preprint arXiv:2406.18665*, 2024.
[7] D. Edge, H. Trinh, N. Cheng, J. Bradley, A. Chao, A. Mody, S. Truitt, and J. Larson, "From local to global: A graph rag approach to query-focused summarization," 2024, microsoft Research, arXiv:2404.16130.
[8] F. Bang *et al.*, "GPTCache: An open-source semantic cache for LLM applications," *GitHub Repository*, 2023, https://github.com/zilliztech/GPTCache.
[9] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom, "Toolformer: Language models can teach themselves to use tools," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
[10] N. Shinn, F. Cassano, A. Gopinath, K. R. Narasimhan, and S. Yao, "Reflexion: Language agents with verbal reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
[11] Y. Sheng, L. Zheng, B. Yuan, Z. Li, M. Ryabinin, B. Chen, P. Liang, C. Ré, I. Stoica, and C. Zhang, "FlexGen: High-throughput generative inference of large language models with a single GPU," *Proc. Int. Conf. Machine Learning (ICML)*, 2023.
[12] LangChain AI, "LangGraph: Building stateful, multi-actor applications with LLMs," https://github.com/langchain-ai/langgraph, 2024.

[13] The Linux Foundation, "Valkey: An open source high-performance key-value store," https://valkey.io/, 2024, fork of Redis OSS 7.2.4.

[14] Chroma, Inc., "ChromaDB: The ai-native open-source embedding database," https://www.trychroma.com/, 2023.

[15] Neo4j, Inc., "Neo4j graph database," https://neo4j.com/, 2024.

[16] Apache Software Foundation, "KRaft: Apache kafka without ZooKeeper," https://kafka.apache.org/documentation/#kraft, 2022.

[17] Prometheus Authors, "Prometheus: Monitoring system and time series database," https://prometheus.io/, 2024.

[18] Grafana Labs, "Grafana alloy: A flexible distribution of OpenTelemetry collector," https://grafana.com/docs/alloy/, 2024.

[19] Anthropic, "Model context protocol specification," https://modelcontextprotocol.io/specification, 2024.

[20] S.-Q. Yan, J.-C. Gu, Y. Zhu, and Z.-H. Ling, "Corrective retrieval augmented generation," 2024. [Online]. Available: https://arxiv.org/abs/2401.15884

[21] B. J. Chan, C.-T. Chen, J.-H. Lai, and D.-B. Wu, "Don't do RAG: When cache-augmented generation is all you need for knowledge tasks," 2024. [Online]. Available: https://arxiv.org/abs/2412.15605

[22] E. Tulving, "Episodic and semantic memory," in *Organization of Memory*. Academic Press, 1972, pp. 381–403.

[23] J. S. Park, J. C. O'Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, "Generative agents: Interactive simulacra of human behavior," 2023. [Online]. Available: https://arxiv.org/abs/2304.03442

[24] C. Packer, S. Wooders, K. Lin, V. Fang, S. G. Patil, I. Stoica, and J. E. Gonzalez, "MemGPT: Towards LLMs as operating systems," 2023. [Online]. Available: https://arxiv.org/abs/2310.08560

[25] M. de Vries, "Paraconsistent logic: A brief introduction," *arXiv preprint*, 2007, arXiv:0707.2161.

[26] K. Gödel, "Zum intuitionistischen aussagenkalkül," *Anzeiger der Akademie der Wissenschaften in Wien*, vol. 69, pp. 65–66, 1932.

[27] J. Łukasiewicz, "O logice trójwartościowej," *Ruch Filozoficzny*, vol. 5, pp. 170–171, 1920.

[28] A. N. Kolmogorov, "Three approaches to the quantitative definition of information," *Problems of Information Transmission*, vol. 1, no. 1, pp. 1–7, 1965.

[29] G. J. Chaitin, "On the length of programs for computing finite binary sequences," *Journal of the ACM*, vol. 13, no. 4, pp. 547–569, 1966.

[30] J. W. Ratcliff and D. E. Metzener, "Pattern matching: The gestalt approach," *Dr. Dobb's Journal*, vol. 13, no. 7, pp. 46–72, 1988.

[31] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen, "A tutorial on thompson sampling," *Foundations and Trends in Machine Learning*, vol. 11, no. 1, pp. 1–96, 2018.

[32] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, and D. Kiela, "Retrieval-augmented generation for knowledge-intensive nlp tasks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[33] V. Krakovna, J. Uesato, V. Mikulik, M. Martic, T. Ord, S. Legg, and J. Leike, "Specification gaming: the flip side of AI ingenuity," DeepMind Blog, 6 April 2020, 2020. [Online]. Available: https://deepmind.google/discover/blog/specification-gaming-the-flip-side-of-ai-ingenuity/