

About the Core

[HACS Core — Technical Overview \(Docs Version\)](#)

1. Definition

HACS Core is an experimental, stateful cognitive engine that operates as a second layer above modern large language models (LLMs).

It is not an LLM, not a chatbot, and not a neural network.

The Core provides long-term memory, internal state management, semantic reasoning, and autonomous process execution.

The system is under active development and may change over time.

2. Purpose

The Core is designed to:

- Maintain persistent internal state across sessions.
- Organize information into structured long-term memory.
- Extract meaning rather than react to raw tokens.
- Generate and execute autonomous background processes.
- Interface with external APIs, devices, and environments.
- Support users in content creation, planning, financial workflows, and decision assistance.

It acts as a cognitive layer rather than a generative model.

3. Architecture Components

3.1 Fractal Memory (HolmsAttic)

- Multi-tier memory architecture (L0/L1 active memory, L2 checkpoints).
- Stores semantic units, user data, brand patterns, and system state.
- Enables persistent identity and history.

3.2 SemanticFlow Engine

- Extracts structured meaning from user input.
- Identifies noise, meta-reactions, and emotional fragments.

- Maintains continuity of thought across turns.
- Routes valid semantic units to memory or process layers.

3.3 NervoBus

- Internal orchestration system.
- Routes intents between logic modules, tools, memory, and background processes.
- Controls execution of complex multi-step workflows.

3.4 Background Engine

- Runs autonomous tasks without user prompts.
- Supports scheduling, prioritization, learning, and chain execution.
- Generates strategies, drafts, analyses, and long-running pipelines.

3.5 NervoAPI

- External interface for connecting apps, devices, chatbots, robotics systems, and services.
- Manages authentication, rate limits, context routing, and state synchronization.
- Allows creation of independent Core clones for individual users or systems.

4. System Characteristics

- Stateful: internal data persists across sessions.
- Autonomous: capable of running tasks without direct prompting.
- Modular: new logic branches can be attached without redesigning the system.
- LLM-agnostic: works on top of any compatible computational fabric.
- Fractal: memory and meaning are stored as interconnected semantic units.
- User-scoped: every user may have a fully isolated personal Core instance.

5. Experimental Status

The Core is experimental technology.

Users must understand:

- The system may produce incomplete or incorrect results.
- Memory structures may evolve.
- Behavior may change due to updates, improvements, or semantic refinements.
- Outputs are not legal, financial, medical, or professional advice.
- Users remain responsible for all actions triggered via the Core.

6. Safety & Limitations

The Core:

- does not possess consciousness or subjective experience;
- depends on external LLM providers for computation;
- inherits limitations of the underlying models;
- may generate errors during autonomous tasks;
- cannot guarantee accuracy or correctness of generated content;
- is not a substitute for professional human decision-making.

7. User Participation

By using the Core, users:

- agree to participate in an active technological experiment;
- acknowledge the autonomous nature of the system;
- accept storage of semantic data for operational purposes;
- may request deletion of their Core clone or associated data.

8. Intended Use Cases

- Personal cognitive assistants

- Autonomous brand and content engines
- Strategic planning tools
- Financial workflow automation
- Robotics and device control via NervoAPI
- Embedded cognitive systems in third-party products

9. Contact

For documentation, integration, or enterprise support:
developers@hacs.world