

EVE-FUSION™ · THE COMPOSITIONAL FABRIC



# Five cooperating reasoning models, composed per request.

*Eve-Fusion™ F5/reasoner — the architecture that powers every product in the MindHYVE family.*

*Eve-Fusion™ is a compositional fabric — not a stack, not a wrapper, not a chain. Five cooperating reasoning models compose dynamically per request, calibrated to the workload, modality, regulatory context, and data-residency constraints. The architecture is plug-and-play and modular: frontier slots are swapped without rebuilding the agent.*

<p>01 <b>Phi-3 Classifier</b></p> <p>Microsoft Phi-3, 3.8B parameters. Sub-200ms routing across workload type, modality, reasoning intensity, and context-length need.</p>	<p>02 <b>Phi-4 Reasoner</b></p> <p>Microsoft Phi-4, LoRA-fine-tuned on the vertical's Eve-Genesis™ edition. The domain-specialized workhorse — carries the majority of workloads.</p>	<p>03 <b>Frontier Slot 01</b></p> <p>Anthropic Claude Opus 4.7. Best-in-class long-form structured reasoning. Used when nuance matters — complex narratives, high-stakes synthesis.</p>
<p>04 <b>Frontier Slot 02</b></p> <p>OpenAI GPT-5.4. Frontier general reasoning, complementary to Opus for cross-validation and ensemble routing. Provider diversity absorbs jurisdictional constraints.</p>	<p>05 <b>Longitudinal Slot</b></p> <p>Meta Llama 4 Scout, 10M-token single-context reasoning. The longest open-weight context publicly available. Full case file, one query.</p>	<p>06 <b>Vision Route</b></p> <p>Meta Llama 4 Vision. Imaging, diagram analysis, document and handwritten review. Same architectural posture as the rest of the stack.</p>

<p><b>&lt; 200ms</b></p> <p>PHI-3 CLASSIFIER ROUTING LATENCY</p>	<p><b>10M</b></p> <p>TOKEN LONGITUDINAL CONTEXT (LLAMA 4 SCOUT)</p>	<p><b>Plug + Play</b></p> <p>FRONTIER SLOTS SWAPPED PER JURISDICTION</p>	<p><b>4</b></p> <p>ONE ARCHITECTURE, FOUR PROVEN VERTICALS</p>
--	---	--	--

## ON THE ARCHITECTURE

Per-product instances: Eve-Healthcare™ for ChironAI, Eve-Education™ for ArthurAI, Eve-Legal™ for JustineAI, Eve-Theology™ for TheoAI. The compound architecture is shared; the Phi-4 reasoner is calibrated to the vertical via LoRA adaptation on the vertical's Eve-Genesis™ synthetic corpus.

## ON TRUST + COMPLIANCE

Provider diversity is structural, not aspirational. Where a jurisdiction or institutional client prohibits a specific provider (Anthropic, OpenAI, Meta), the relevant slot is swapped without rebuilding the agent or retraining the reasoner. The architecture survives regulatory change.

CONTINUE READING → [02 · MindHYVE](#)    [10 · Eve-Genesis](#)    [11 · Differentiation](#)    → [/tech](#)

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries  
[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

[mindhyve.ai](https://mindhyve.ai)