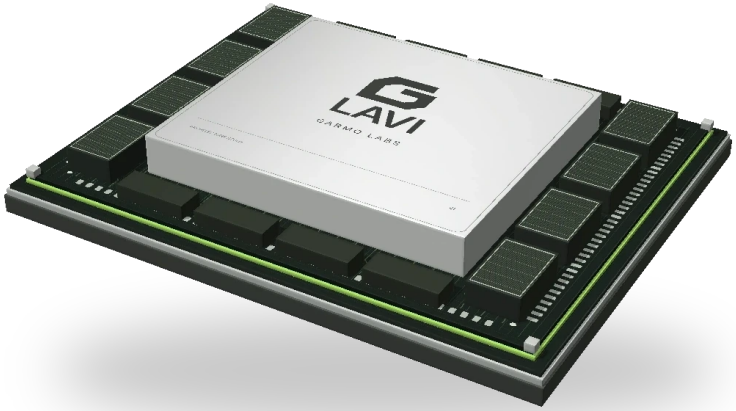
MODEL-NATIVE AI

LAVI

A frontier AI architecture that brings the model, its memory and its execution machinery into one negotiated system.



Architectural concept. Not a die photograph, physical layout or measured device.

LAVI · Integrated reference · Architecture + RTL

Function and architecture

- Model–silicon execution contracts
- Tensor, attention, recurrent and MoE engines
- Semantic KV and persistent execution graphs
- Scaled LAVI-S130 proof-chip path

Integration summary

Model DNA and Silicon DNA negotiate an execution contract. Compute, memory-intelligence and fabric partitions compose the larger architecture; LAVI-S130 is the scaled proof vehicle.

Technical record

Canonical identity	LAVI
Responsibility	A frontier AI architecture that brings the model, its memory and its execution machinery into one negotiated system.
Delivery boundary	Architecture + RTL
Evidence scope	Integrated reference

Delivery and release binding

- Exact SoC/top and configuration
- Clock, reset, power and pad contract
- Memory, PHY and package dependencies
- Release-specific PVT and signoff views

Confirm for the selected release: Node/library binding, Clock/reset limits, Physical and protocol dependencies, Artifact hashes, Acceptance criteria.

Evidence and qualification

Our records show integrated chip/model semantic checks and a 150,000-check LAVI-S130 reference campaign. Logical chip-top, proof-core and OpenFrame wrapper source are present.

Reference execution is not measured frontier silicon. Final HDL/physical qualification and trained frontier-scale checkpoints remain separate gates.

Evaluation and licensing

Evaluation: Architecture and release evaluation. Production: Chip / SoC licensing program. Custom work: Custom silicon integration. Rights, pricing, support and acceptance criteria are agreed for the exact scope.

Contact: patrick@garmolabs.com | Product path: /silicon/lavi