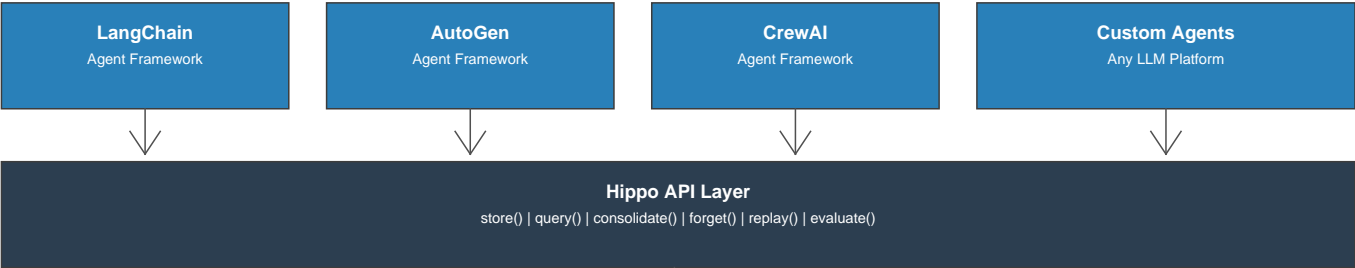
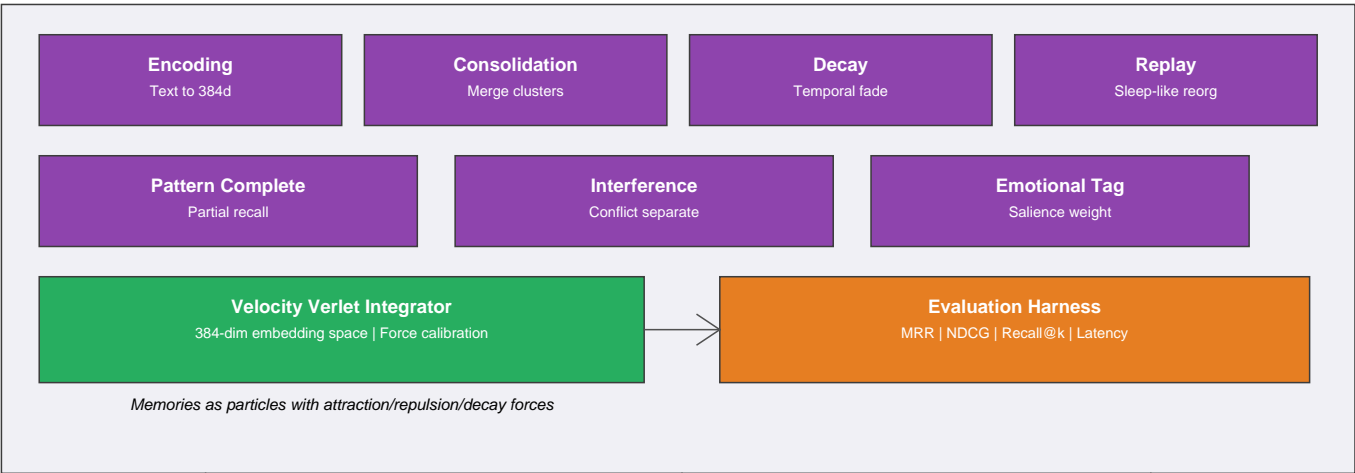


Page 1: System Architecture

AGENT LAYER



HIPPO CORE ENGINE



PERSISTENCE LAYER



Current System Status (v0.27.0)

28 source modules | 527 passing tests | 6 platform integrations | 384-dim embeddings | 7 hippocampal mechanisms | MIT + proprietary extensions

Phase 2 Vision: Multi-Agent Shared Memory

Multiple foundation model agents contributing to a collective particle space with emergent knowledge structures

Page 2: Validation and Benchmark Matrix

A. Benchmark Comparison Matrix

Metric	Hippo	FAISS	ChromaDB	LlamaIndex	Success Criterion
MRR (1K entries)	Measured	Measured	Measured	Measured	> 0.70
MRR (10K entries)	Measured	Measured	Measured	Measured	> 0.70
MRR (50K entries)	Measured	Measured	Measured	Measured	> 0.70
MRR (100K entries)	Measured	Measured	Measured	Measured	> 0.70
NDCG@10 (1K)	Measured	Measured	Measured	Measured	> 0.75
NDCG@10 (100K)	Measured	Measured	Measured	Measured	> 0.75
Recall@5 (100K)	Measured	Measured	Measured	Measured	> 0.80
Recall@20 (100K)	Measured	Measured	Measured	Measured	> 0.90
Retrieval latency (1K)	Measured	Measured	Measured	Measured	< 50ms
Retrieval latency (100K)	Measured	Measured	Measured	Measured	< 500ms
Precision after 30 days	Measured	Measured	Measured	Measured	< 5% degradation
Memory consolidation	Yes	No	No	No	Unique capability
Graceful forgetting	Yes	No	No	No	Unique capability
Sleep-like replay	Yes	No	No	No	Unique capability

B. Validation Plan Matrix

Obj.	What We Validate	Method	Success Criterion	Deliverable
O1	Convergence of particle dynamics under load	Lyapunov stability analysis + 100hr empirical tests	Bounded energy; convergence in all 10 workload profiles	Convergence proof document
O2	Retrieval quality vs state-of-the-art RAG	Paired benchmarks vs FAISS, ChromaDB, LlamaIndex (30 runs each)	MRR > 0.70, NDCG > 0.75 at 100K entries; < 5% precision degradation	Benchmark report with statistical analysis
O3	Multi-agent shared memory feasibility	Architecture analysis; conflict resolution simulation	Specification complete; 2+ enterprise partners engaged	Phase 2 technical specification

C. Feasibility Study Timeline

Phase	Month 1 (Oct 2026)	Month 2 (Nov 2026)	Month 3 (Dec 2026)
O1: Convergence			
O2: Benchmarks			
O3: Multi-agent spec			
Reporting			
Partner engagement			