

MyAi AegisWeave: A Zero-Trust, Memory-Centric Architecture for Private Multi-Agent LLM Systems

Le Grande Maison de Beausoleil

MyAI / iownmyai.com

<https://iownmyai.com>

July 2026

Abstract

MyAi AegisWeave is the architectural fabric behind the MyAi ecosystem: a zero-trust, memory-centric, multi-agent large language model (LLM) system designed to run on privately owned hardware. It is optimized for high-net-worth estates, enterprises, and government environments that require strong security, strict privacy, and efficient edge operation. This whitepaper introduces the AegisWeave design principles, layered architecture, memory and storage model, agent taxonomy, and domain modules for climate, security, energy, entertainment, appliances, research, and private R&D.

Contents

1	Introduction	4
1.1	Motivation	4
1.2	Design Goals	4
1.3	Scope of This Document	4
2	High-Level Architecture	4
2.1	Layered System Overview	4
2.2	Physical and Energy Layer	4
2.3	Memory and Storage Layer	4
2.4	Model and Inference Layer	4
2.5	Agent and Orchestration Layer	4
2.6	Security and Governance Layer	4
2.7	Experience and Integration Layer	4
3	Security and Privacy Doctrine	4
3.1	Zero-Trust Principles	4
3.2	Data Residency and Privacy	4
3.3	Model and Agent Security	4
3.4	Auditability and Governance	4
4	Memory-Centric Design	4
4.1	Rationale for Memory-Centric AI	4
4.2	Memory and Storage Tiers	4
4.3	Micron-Based Configurations	5
4.4	Ubiquiti Sensing and Local Storage	5
5	Model and Inference Core	5
5.1	Transformer + Mixture-of-Experts Backbone	5
5.2	Multi-Head Latent Attention and KV Compression	5
5.3	Quantization and Edge Constraints	5
5.4	Inference Topologies	5
6	Agent Architecture	5
6.1	Agent Taxonomy	5
6.2	Agent Communication and Message Bus	6
6.3	Policy and Guardrail Enforcement	6
7	Domain Modules	6
7.1	Climate and HVAC Agent	6
7.2	Lighting and Ambience Agent	6
7.3	Security and Defense Agent	6
7.4	Energy Agent (MyAi Energy Vault)	6
7.5	Entertainment and Media Agent	6
7.6	Appliances Agent	6
7.7	Research and Education Agent	6
7.8	Creative Agent	6

7.9	Technical Disciplines Agents	6
7.9.1	Physics / Engineering / Mathematics	6
7.9.2	Astronomy	6
7.10	Economics and Trading Agent	6
7.11	Jurisprudence and Governance Agent	6
8	Internal Climate, Sensing, and Control Integration	6
8.1	Device and Bus Integration	6
8.2	Digital Twin Model	6
8.3	Safety, Fail-Safe, and Degradation Modes	6
9	Implementation Considerations	6
9.1	Deployment Topologies	6
9.2	Performance and Capacity Planning	6
9.3	Operational Playbooks	6
10	Conclusion and Roadmap	6
10.1	Summary of AegisWeave Principles	6
10.2	Future Work	6

1 Introduction

1.1 Motivation

1.2 Design Goals

- Strong security and zero-trust by default.
- Privacy-first, local-by-default operation.
- Efficiency via sparse computation and memory-centric design.
- Modular expansion from a single node to full estates and enterprises.

1.3 Scope of This Document

2 High-Level Architecture

2.1 Layered System Overview

Figure 1: Conceptual layered architecture of MyAi AegisWeave (illustration placeholder).

2.2 Physical and Energy Layer

2.3 Memory and Storage Layer

2.4 Model and Inference Layer

2.5 Agent and Orchestration Layer

2.6 Security and Governance Layer

2.7 Experience and Integration Layer

3 Security and Privacy Doctrine

3.1 Zero-Trust Principles

3.2 Data Residency and Privacy

3.3 Model and Agent Security

3.4 Auditability and Governance

4 Memory-Centric Design

4.1 Rationale for Memory-Centric AI

4.2 Memory and Storage Tiers

- DRAM/HBM: fast, small, for inference KV and hot state.
- NVMe SSD (Tier 1): model weights, vector indices, hot RAG corpora.
- NVMe SSD / NVR (Tier 2): cold R&D data, surveillance archives, logs.

Figure 2: Indicative memory hierarchy for AegisWeave (placeholder).

4.3 Micron-Based Configurations

4.4 Ubiquiti Sensing and Local Storage

5 Model and Inference Core

5.1 Transformer + Mixture-of-Experts Backbone

5.2 Multi-Head Latent Attention and KV Compression

5.3 Quantization and Edge Constraints

5.4 Inference Topologies

6 Agent Architecture

6.1 Agent Taxonomy

Table 1: Agent taxonomy: core, domain, and tool agents (placeholder).

Category	Agent Type	Example Roles
Core	Conversation / Avatar	Primary user interface, session memory
	Orchestration	Routing, task decomposition
	Security / Policy	AuthZ, redaction, approvals
	Health / Telemetry	Monitoring, anomaly detection
Domain	HVAC / Climate	Comfort, energy-optimal control
	Lighting / Ambience	Scenes, circadian lighting
	Security / Defense	Cameras, K9, drones, access
	Energy	Energy Vault, solar, storage, grid
	Entertainment	Media, scenes, events
	Appliances	Scheduling, efficiency
	Research / Education	Tutoring, corpus management
	Creative	Video, audio, image, document creation
	Physics / Eng / Math	Technical reasoning, modeling
	Astronomy	Observations, event tracking
	Economics / Trading	Portfolio, scenarios, execution
	Jurisprudence	Contracts, policy, governance
Tool	FS / KB Access	Files, knowledge store, vector DB
	Simulation / Monte Carlo	Stochastic analysis
	Viz / Plotting	Charts, diagrams
	Secure Network	Trading APIs, external LLMs (gated)
	Device / Bus	Modbus, BACnet, KNX, CAN, etc.

6.2 Agent Communication and Message Bus

6.3 Policy and Guardrail Enforcement

7 Domain Modules

7.1 Climate and HVAC Agent

7.2 Lighting and Ambience Agent

7.3 Security and Defense Agent

7.4 Energy Agent (MyAi Energy Vault)

7.5 Entertainment and Media Agent

7.6 Appliances Agent

7.7 Research and Education Agent

7.8 Creative Agent

7.9 Technical Disciplines Agents

7.9.1 Physics / Engineering / Mathematics

7.9.2 Astronomy

7.10 Economics and Trading Agent

7.11 Jurisprudence and Governance Agent

8 Internal Climate, Sensing, and Control Integration

8.1 Device and Bus Integration

8.2 Digital Twin Model

8.3 Safety, Fail-Safe, and Degradation Modes

9 Implementation Considerations

9.1 Deployment Topologies

9.2 Performance and Capacity Planning

9.3 Operational Playbooks

10 Conclusion and Roadmap

10.1 Summary of AegisWeave Principles

10.2 Future Work