AIGO Architecture

AI Governance Orchestration Framework

High-level Architecture & Orchestration Capability Overview

1. Architectural Overview

AI Governance Orchestration and NLP framework is designed for organisations that prioritise data privacy, operational safety, and trust - whether operating in regulated industries or simply handling sensitive information. Unlike typical agent frameworks that focus solely on runtime execution, AIGO enforces predefined capability contracts, execution boundaries, and governance filters at the infrastructure level, ensuring security, auditability, and adaptability from day one.

2. Core Architectural Components

2.1 Agent Orchestration Layer

- **Execution Model:**
  Multi-agent environment where each agent's scope, allowed actions, dataset access, and tool/API access are defined before runtime in machine-readable capability contracts.
- **Delegation Governance:**
  - **Role-scoped routing rules** ensure only authorised agents handle specific tasks.
  - **Delegation depth limits (chained tasks and handovers)** prevent infinite or uncontrolled nested calls, ensuring bounded execution.
  - **Typed function delegation** ensures outputs match the expected structure before downstream execution.
- **Compliance Integration:**
  Orchestration rules can enforce language, currency, and content filters inline with the task execution.

2.2 Tool & API Integration Layer

- **Dual Integration Model:**
  - **APIs** - External tools exposed by the client's environment.
  - **Tools** - Internal APIs exposed by AIGO for orchestration.
- **Shared Functionality:** Both APIs and Tools share:
  - Visibility flags (`DisplayToModel`, `DisplayToUser`, `SaveToMemory`)
  - Parameter definitions, validation rules, and allowed caller lists.
  - API-sourced control - inclusion of RAG, request failure or success recommendations.
- **Ownership Separation:**
  - **API Configs** - Managed by clients, reflecting their external integrations.

    - **Tool Configs** - Managed internally, forming AIGO's internal orchestration toolbox.
- **Security & Compliance:**
  Configs dictate access scope, approved parameter types, and result handling before any execution.

2.3 Data Preprocessing & Schema Handling

- **Schema-Agnostic Processing:**
  Automatically interprets and processes datasets with varying schemas, avoiding brittle, hard-coded mappings.
- **Domain-Specific Optimisation:**
  Summarisation tailored to regulated domains - ensures outputs are AI-usable but free of sensitive identifiers.
- **Dataset Lineage:**
  Every processed dataset is tied to a **persistent dataset ID**, enabling full traceability from ingestion to output.

2.4 Memory Management & State Control

- **Scoped Memory Model:**
  Key-value structured memory with configurable persistence rules ensures context is retained only when appropriate.
- **Memory Types:**
    - Ephemeral - duration of the workflow
    - Session team-level - accessible by agents within an identified team
    - Session agent-level - accessible by individual agents
    - Global - accessible across sessions
- **Automatic Summarisation:**
  Older or bulkier memory entries are summarised according to pre-set rules, preventing context bloat.
- **Isolation Guarantees:**
  Prevents leakage of sensitive data between unrelated workflows or agents.
- **Full Session Capture:**
  Logs every message, inter-agent exchange, and tool/API call - even those hidden from the end user.

2.5 Compliance & Audit Framework

- **Inline Output Governance:**
  Automatically enforces formatting, currency, language, and content constraints before returning outputs.
- **Immutable-Ledger Ready:**
  All logs can be committed to blockchain storage for tamper-proof audit trails.
- **Per-Client Storage Isolation:**
  Each client's logs, configs, and datasets are stored in dedicated containers for compliance and data sovereignty.

2.6 Extensibility & Scaling

- **Config-Driven Expansion:**
  New tools, APIs, and agents can be introduced without code changes - configuration is the primary onboarding mechanism.
- **Cloud-Native Horizontal Scaling:**
  Architecture supports load balancing across multiple orchestration instances with every internal tool deployed on a micro-service, including execution of queries, data persistence, analysis scripts building etc.
- **Multi-Tenant with Isolation:**
  Designed for SaaS deployment where tenants share infrastructure but remain logically and operationally isolated.

3. Current Capability vs Roadmap Fit

- **Today:**
  - Fully capable of NLP Q&A against structured and semi-structured datasets.
  - Multi-agent orchestration is already functional with compliance controls in place.

- **Future-Ready:**
  - Same orchestration backbone can support workflow automation, hybrid human-AI decisioning, and adaptive routing without redesign.
  - Native integration model allows gradual layering of additional tools and external APIs as the use case evolves.

4. AIGO vs Traditional MCP

Capability	AIGO	Traditional MCP
Execution Control	Pre-runtime capability contracts, config-driven routing	Runtime host-driven
Delegation Depth Control	Yes (per-agent or global)	No native feature
Output Governance	Built-in compliance enforcement	Manual or external
Data Lineage & Auditability	Persistent dataset IDs, full logging	Minimal to none
Schema Flexibility	Schema-agnostic preprocessing	Typically schema-locked
Visibility Rules	Per-tool/API control for user/model/memory	Limited or absent
Integration Model	Dual: client APIs + internal tools	Single registry
Memory Control	Scoped, summarised, isolated	Often uncontrolled
Blockchain Audit Option	Yes (optional)	Not standard

As of March 2025

5. Why This Architecture is Future-Proof

- **NLP-first, Orchestration-ready:** No architectural rewrite needed to progress from Q&A to orchestrated workflows.
- **Compliance by Design:** Essential for scaling in regulated markets - prevents costly retrofits.
- **Config-driven Extensibility:** Reduces engineering overhead for expansion and client-specific adaptations.
- **Defensible IP Features:** Unique governance and orchestration model not found in generic MCP implementations.

Notice

This document contains proprietary information of **One Horizon Technology** relating to its technology, architecture, and intellectual property. It is provided solely for the purpose of evaluation by the recipient and may not be copied, disclosed, or used for any other purpose without the prior written consent of Areto Net Limited. Any unauthorised use, reproduction, or distribution of this material, in whole or in part, is strictly prohibited.

Primary contact: hello@onehorizon.tech