ANASSA

Autonomous Neural Agents for Spatial Systems Architecture

Spatial Agentic AI Orchestration Framework

A new approach to spatial agentic AI orchestration

SHORT WHITE PAPER

June 2026

Website: anassa.ai **LinkedIn:** linkedin.com/company/anassa-ai **Email:** contact@anassa.ai

Overview

ANASSA is a **Spatial Agentic AI Orchestration Framework**. A master orchestrator, the **Queen**, decomposes intent into a coordinated task graph, sequences specialized AI agents through reasoning and execution, validates every output against authoritative ground truth, and synthesizes the result into trusted, auditable intelligence. ANASSA sits **above** existing data systems, pipelines, and compute rather than replacing them, and it treats Earth and Space as one unified spatial-intelligence domain.

Status

A full applied research paper is currently under review at a leading AIS (Association for Information Systems) conference. The next step is implementation-level validation and benchmarking, taking the framework from architecture to production-grade infrastructure.

Author: **Constantinos Papantoniou, Ph.D.**, Execution Architect and Founder

Acknowledgment: **Brian Hilton, Ph.D.**, Former Clinical Full Professor of IS&T

1. Introduction

Spatial intelligence has never been more abundant, or more fragmented. Earth-observation satellites, orbital sensors, ground networks, and field systems produce vast streams of geospatial data, yet the value locked inside that data is increasingly limited not by the data itself but by the human effort required to coordinate it across disconnected systems [4].

This white paper introduces **ANASSA**, a Spatial Agentic AI Orchestration Framework that closes that coordination gap by placing specialized AI agents under a single orchestrator that senses, reasons, plans, validates, and acts across spatial data, under human governance. The framework builds on prior research into spatially-enabled distributed and trusted systems [1], [10], [11], and on the rapid maturation of large language models and autonomous AI agents [2], [3], [5], [6].

This document is an overview of the framework: its motivation, architecture, governance model, and four applied use cases across Earth and Space. It is intended for technical and strategic readers evaluating how agentic AI can be applied to high-stakes spatial decision-making.

2. The Problem: The Fragmentation Tax

Modern geospatial operations face a coordination crisis. Across governments, enterprises, defense, and research, spatial data flows through disconnected systems: GIS platforms, satellite feeds, IoT sensors, field devices, cloud databases, and legacy infrastructure, each operating in isolation. Humans are forced to act as the connective tissue, manually extracting, transforming, routing, and reconciling data across these fragmented systems.

This coordination burden has a name: the **Fragmentation Tax**. It is the measurable loss in capacity that occurs when human time and cognitive bandwidth are consumed by integration work rather than spatial reasoning and decision-making. Every hour an analyst spends stitching pipelines together is an hour not spent answering the spatial question that drives the decision.

> **Central Research Premise**
> *What if spatial decisions took seconds, not hours? What if Earth and Space data could coordinate themselves?*

3. What ANASSA Is

ANASSA (**A**utonomous **N**eural **A**gents for **S**patial **S**ystems **A**rchitecture) is a research architecture, not a single software product. It defines how distributed geospatial AI agents can autonomously sense, reason, plan, decide, execute, and learn across industry sectors, data sources, deployment environments, and spatial domains [2], [3]. The shift it researches is from **Systems of Insight**, maps and dashboards that help humans understand, to **Systems of Action** that close the loop between spatial data and spatial decisions at machine speed, under human governance.

Core Design Principles

- **Agentic by design.** Every functional layer is built around autonomous agents, not passive tools.
- **Centrally coordinated.** All agents operate under one master orchestrator that manages tasking, communication, and alignment.
- **Human-in-the-loop at critical points.** Mandatory review gates sit at Planning and Decision, where domain judgment and accountability matter most.

- **Trusted and auditable.** Every decision, reasoning chain, and output is logged, explainable, and checked against ground truth.
- **Continuous learning.** Each completed cycle feeds improvements back into every component.
- **Deployment-agnostic.** Supports both large enterprise deployments and lightweight implementations.

4. How It Works: The Cognitive Loop

At the center of ANASSA is a continuous, closed feedback loop coordinated by the Queen orchestrator. The orchestrator is not a router. It is the **Execution Architect** of the system: it decomposes intent into agent tasks, sequences execution across the framework, manages inter-agent communication, enforces governance, and synthesizes results. It also absorbs the coordination and reporting work that normally falls on people, so the loop runs with minimal manual overhead.

S1 Sense	**Sensing Agent.** Detects spatial events and anomalies in incoming data that warrant processing.
S2 Reason	**Reasoning Agent.** Applies spatial intelligence to understand context and patterns; selects models, tools, and workflows.
S3 Plan	**Planning Agent.** Generates workflows and solution strategies aligned with participant needs and industry rules. **Human-in-the-loop**
S4 Decide	**Decision + Validation.** Evaluates options and validates all spatial claims against authoritative ground truth. **Human-in-the-loop**
S5 Execute	**Execution Agent.** Deploys solutions, publishes maps and dashboards, and triggers downstream workflows.
S6 Learn	**Learning Agent.** Captures outcomes and feeds improvements back to every component. Each cycle improves on the last.

Two **mandatory human-in-the-loop gates** sit at Planning (S3) and Decision validation (S4), keeping the highest-stakes choices under explicit human accountability. The shift is not from humans to machines. It is from humans stitching systems together to humans governing intelligent systems that act.

5. Framework Architecture

ANASSA organizes **47 specialized agents** into eleven numbered components across four layers, all under one orchestrator. A governance overlay and a trusted data foundation bound a central cognitive engine, so every output is auditable, reproducible, and grounded in authoritative data, the properties an infrastructure platform needs to scale safely.

- **Governance & Output.** Governs all agent and workflow behavior and certifies the framework's trusted intelligence outputs.
- **Operational.** Participants, industry sectors, deployment paths, the system-of-systems backbone, and multimodal spatial data.
- **Cognitive Engine.** The six-step Cognitive Loop, the reasoning and action core described above.
- **Foundation.** Trusted authoritative data, the knowledge base, and the process-flow connections linking all components.

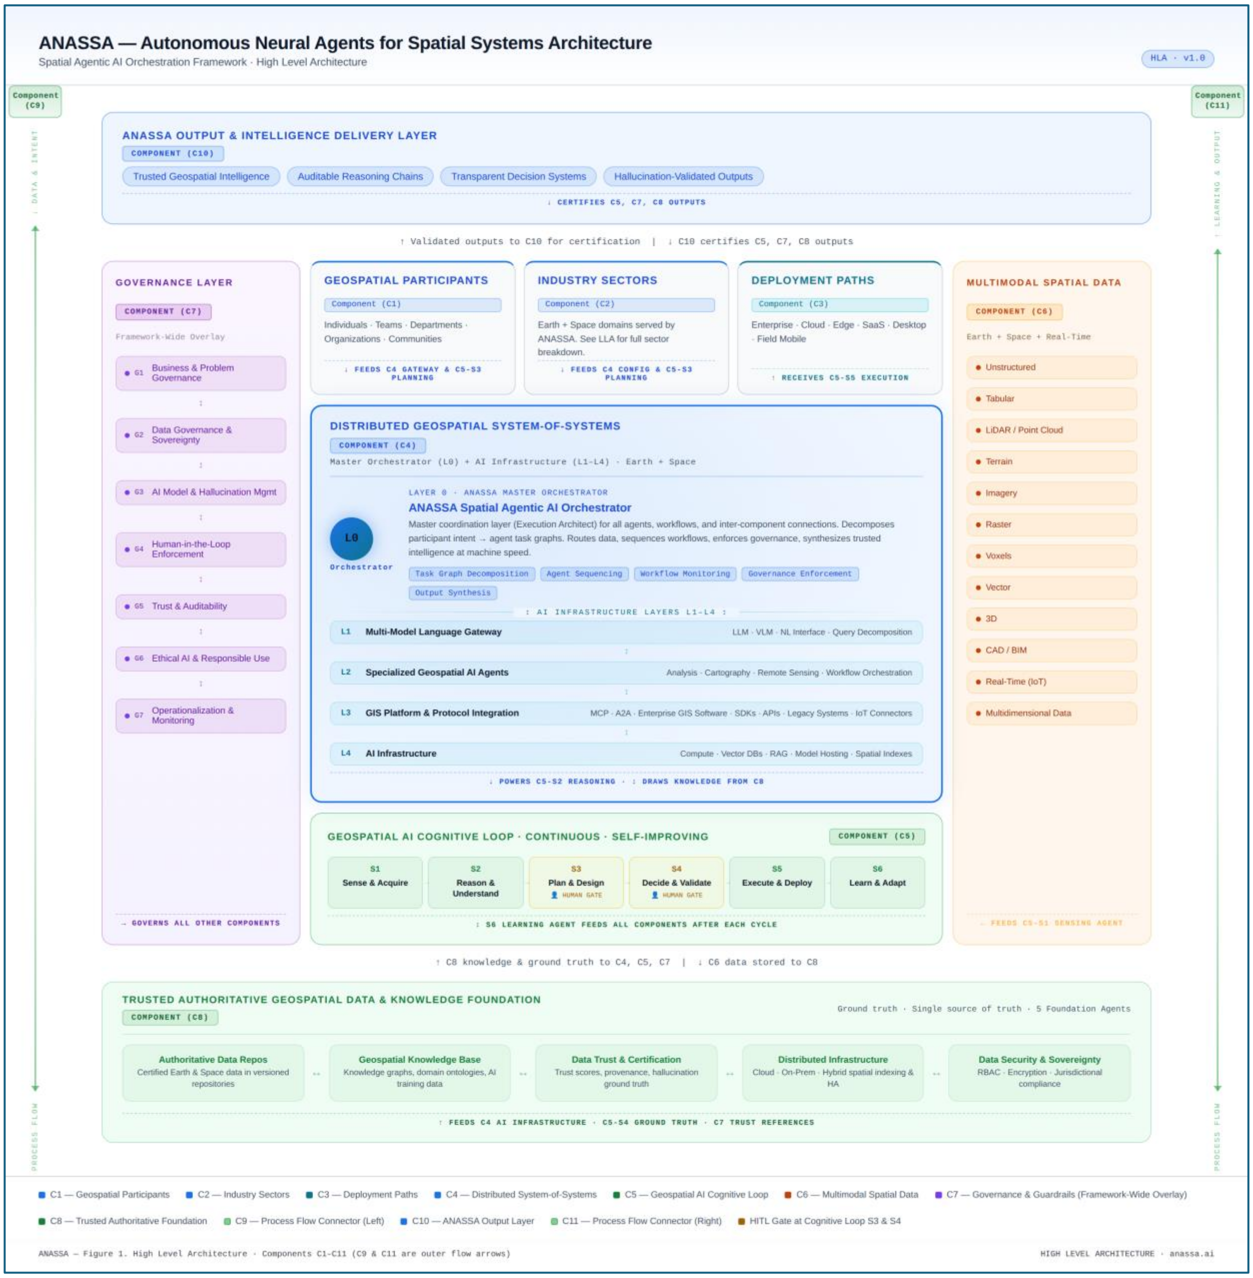

Figure 1. ANASSA High-Level Architecture: governance overlay, the orchestrator, the four-layer system-of-systems, the six-step cognitive loop with human-in-the-loop gates, multimodal data, and the trusted data foundation.

6. Governance, Trust, and Hallucination Management

In the spatial domain, a hallucinated output is not merely a wrong answer. A fabricated coordinate, an invented boundary, or a false spatial claim can produce an incorrect map used for a high-stakes decision, drive wrong infrastructure choices, create jurisdictional disputes, and, in safety-critical contexts, endanger lives. ANASSA therefore treats **hallucination management as a first-class governance responsibility**: every numerical and spatial claim is validated against primary source data before it is allowed to influence a decision.

Governance is a framework-wide overlay informed by the CPMAI lifecycle for AI projects [7]. It spans problem definition, data governance and sovereignty, model and agent behavior, human oversight, full audit logging, ethical and responsible use, and continuous monitoring. Trust is not assumed; it is verified and recorded.

7. Four Applied Use Cases

Two enterprise data cases and two high-stakes, machine-pace cases, spanning Earth and Space. Each takes existing data and compute and tests what the orchestration layer adds in reliability, speed, and scale. These are research proposals; no proof-of-concept has yet been run.

01 Cross-Source Data Reconciliation EARTH

What. Reconciling ownership and boundary records across overlapping, inconsistent datasets, a multi-step workflow where automated pipelines often fail mid-execution on schema, projection, or parameter mismatches.

How ANASSA helps. ANASSA detects execution failures, reconfigures the workflow through its feedback loop, and validates results against authoritative source data, producing a reliable, auditable pipeline rather than a brittle one-shot run.

02 Supply Chain Risk EARTH

What. Assessing supply-chain vulnerability across warehouse locations, transport routes, and environmental hazards such as flood zones, a workflow that breaks when intermediate outputs are missing.

How ANASSA helps. The orchestrator monitors progress across steps, detects inconsistencies such as missing layers, and replans on the fly, while the sensing layer ingests live hazard feeds so the analysis adapts as conditions change.

03 Autonomous Vehicle Safety Simulation EARTH

What. AV systems are validated through simulation, and the highest-value runs are rare, safety-critical scenarios [8]. The constraint is not raw compute but targeting: uniform simulation spends the same effort on a straight highway as on a complex intersection.

How ANASSA helps. ANASSA deploys its cognitive loop above the simulation pipeline. It flags zones where real-world risk exceeds current coverage, produces a ranked priority plan reviewed and approved by engineers, injects approved seeds into the pipeline, and measures coverage efficiency per unit of compute each cycle.

04 Lunar Infrastructure Site Selection SPACE

What. Selecting permanent lunar infrastructure sites means fusing many independent domains: solar power, thermal, radiation, terrain, communications geometry, and water-ice proximity, each from a different orbital instrument at a different resolution and reference frame [9], [12].

How ANASSA helps. ANASSA co-registers heterogeneous datasets into one reference frame, scores candidates in parallel across every domain, validates numbers against primary source data, and compresses multi-month specialist review into a structured, reusable, human-approved package.

8. Conclusion

Across all four cases the common thread is the same: **sequencing compute, managing dependencies, ensuring reproducibility, keeping humans in the loop, and treating reliability as a safety commitment**. ANASSA is fundamentally an execution-architecture challenge: the discipline of how distributed agents, compute, and data are sequenced and governed. The orchestration approach generalizes across domains, which is the point: it is an architecture for coordinating distributed spatial systems, not a single-purpose tool. The data exists and the AI foundations exist; what has been missing is the orchestration layer that connects them into a coherent, trusted, and auditable intelligence system.

References

[1] Papantoniou, C. (2021). GeoBlockchain: The Analysis, Design, and Evaluation of a Spatially Enabled Blockchain. Ph.D. Dissertation, Claremont Graduate University, School of Information Systems and Technology. CGU Theses and Dissertations, 284. https://scholarship.claremont.edu/cgu_etd/284

[2] Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W., Wei, Z., and Wen, J. (2024). A Survey on Large Language Model based Autonomous Agents. Frontiers of Computer Science, 18(6), 186345. https://doi.org/10.1007/s11704-024-40231-1

[3] Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., et al. (2023). The Rise and Potential of Large Language Model Based Agents: A Survey. arXiv preprint arXiv:2309.07864.

[4] Blaschke, T., Hay, G.J., Kelly, M., Lang, S., Hofmann, P., Addink, E., Feitosa, R.Q., van der Meer, F., van der Werff, H., van Coillie, F., and Tiede, D. (2014). Geographic Object-Based Image Analysis: Towards a New Paradigm. ISPRS Journal of Photogrammetry and Remote Sensing, 87, 180-191.

[5] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., et al. (2020). Language Models are Few-Shot Learners. Advances in Neural Information Processing Systems, 33, 1877-1901.

[6] Anthropic. (2024). The Claude Model Family. Technical Report. https://www.anthropic.com/research

[7] Cognilytica. (2019). CPMAI: Cognitive Project Management for AI Methodology. Cognilytica Research. https://www.cognilytica.com/cpmai-methodology/

[8] Waymo. (2024). Behind the Innovation: AI and ML at Waymo. Waymo Engineering Blog, October 2024. https://waymo.com/blog/2024/10/ai-and-ml-at-waymo/

[9] Heiken, G., Vaniman, D., and French, B.M. (Eds.). (1991). Lunar Sourcebook: A User's Guide to the Moon. Cambridge University Press.

[10] Papantoniou, C. (2024). Trusting Spatial AI in the Battlespace. LinkedIn Pulse. https://www.linkedin.com/pulse/trusting-spatial-ai-battlespace-papantoniou-ph-d-mba-gisp--nxf3c/

[11] Papantoniou, C. (2024). Why Spatial AI Agents Need Domain Expertise to Succeed. LinkedIn Pulse. https://www.linkedin.com/pulse/why-spatial-ai-agents-need-domain-expertise-succeed-constantinos-bcslc/

[12] NASA. (2023). Lunar Reconnaissance Orbiter: Science Instruments and Data Products. NASA Goddard Space Flight Center. https://lro.gsfc.nasa.gov/