SEAGLASS FOUNDRY

Forged For Speed. Refined by Craft.

DEFENSE & INTELLIGENCE

ForgeMind™

Trustworthy Natural-Language Geospatial for Mission Environments

A natural-language control plane for GPU geospatial analysis that refuses rather than fabricates. ForgeMind™ turns an operator's spoken or typed request into a validated ForgeGIS pipeline, and when a request cannot be honored correctly it returns an explicit, machine-readable refusal — never a plausible-but-wrong product. Across a 6,696-turn validation campaign it produced zero confident-wrong results. Pure-JVM, local-disk, LLM-agnostic, and deployable where the mission runs.

Zero

confident-wrong results

79 conversations, 6,696 turns; June 2026 validation campaign

LLM

agnostic

Claude or OpenAI behind one abstraction; self-hosted models the same integration point

Pure

JVM, local-disk

no required external service beyond the model endpoint

Refusal by construction

The language model is confined to understanding intent. Deterministic machinery validates every pipeline against a type system before any work runs. A request that cannot be mapped to a valid operation, lacks its parameters, or cannot bind to data is refused with a specific reason — never run speculatively.

Provider independence

ForgeMind runs against any supported model behind one abstraction. Switch providers to follow availability, cost, or procurement constraints — including continuity during a provider outage — with a configuration change, not a code change. For programs hosting their own models, that abstraction is the integration point.

Determinism inherited from ForgeGIS

Numerical results come from ForgeGIS and are reproduced faithfully because ForgeMind dispatches to the same compute implementations whether an analyst or an agent makes the call. The audit and reproducibility guarantees regulated work requires hold across the natural-language surface.

CONTINUED ON REVERSE →

VALIDATED WHERE IT COUNTS

A campaign built around the cases a careless system gets wrong

The validation campaign exercises ForgeMind across realistic multi-turn operator sessions — establishing a subject, refining it across turns, chaining dataset discovery into analysis, and deliberately probing the negative cases. Unresolved references (a name never defined, an ambiguous "the other one," a follow-up with no antecedent) are refused with a specific reason rather than guessed at. The headline result is the absence of a failure category: zero confident-wrong answers across the run.

The line-of-sight surface, reachable by intent

Viewshed, multi-observer viewshed, intervisibility, and terrain-constrained suitability — the propagation and visibility workflows defense and IC depend on — are reachable in natural language and composable into multi-step pipelines on the GPU-resident ForgeGIS engine. Named map entities (an objective, a route, a sensor site) are referenced by stable handle and resolved against real geometry, never re-derived from a bounding box.

PROOF POINTS

0
confident-wrong results across the entire validation campaign

97
voice-routable capabilities exercised end to end

~25×
lower LLM cost than an agent-loop approach to the same workflows

> **Fail-closed on geometry.** When the true geometry of a referenced entity is not reachable, the operation defers and asks rather than emitting a plausible-but-wrong product — the same refusal discipline that governs the rest of the system, applied to spatial references. Before this discipline, "buffer the channel" could silently produce one circle at the centroid of a bounding box. It no longer can.
>
> **Deploys where the mission runs.** ForgeMind is a pure-JVM application with local-disk persistence; its compute, catalog, and map-client backends run as local subprocesses over stdio rather than the network, and its memory store is a single local database file. The only external dependency is the model endpoint — and for programs hosting their own models, that is one configuration line.

Request the Technical Brief and a demonstration

rich@seaglassfoundry.com · seaglassfoundry.com

The Technical Brief covers the architecture, the validation campaigns, and deployment guidance. An evaluation build is available to qualified programs on request.